

Watson & Walker zRoadshow Winter 2023 Edition

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



Session 94032 – March 2023

Introduction



- Who are we?
- Watson & Walker founded in 1988 by **Cheryl Watson & Tom Walker**.
- Publisher of [Cheryl Watson's Tuning Letter](#) and *Cheryl's CPU Charts* since 1991.
- We are completely independent, not beholden to any vendor.
- Our focus areas at the moment are helping clients prepare for z16, and convincing senior management that the zSkills issue is not fake news.
- For more information, see our website, www.watsonwalker.com.
- The objective of *this* session is to make you aware of interesting tidbits that we have come across recently, and point you to where you can get more information.

Highlights of the Week

- **Cheryl's:**

- **Kurt's** [A Tour of z/OSMF Software Management](#) session – “unless you’re planning to be off z/OS within a year, you’re insane not to be using z/OSMF in production for ALL its applications that apply to your shop”
- [What's New in DFSMS](#), **Barb McDonald**
 - z/OS 3.1 – Cloud Data Access (CDA) Method; move to cloud bypassing tape and disk
 - z/OS 2.4/2.5 w/APARs – GDKUTIL to transfer between z/OS and cloud
 - z/OS 3.1 – Adds STCKE to catalog records to allow better catalog recovery!!!!!! (One of Cheryl's oldest SHARE requirements)
- [MVS Program Opening - z/OS Service and APARs](#), **John C. Shebey III**
 - JES2 2.5 APAR OA63943, checkpoints on CF, PTFs on 2/22.
 - RSM 2.4/2.5 Open APAR – SDUMP option with OPTIMIZE=YES missing 64-bit virtual memory object (affects CICS and OMVS).
 - IBM z/OS Academy – <https://ibm.biz/zOSAcademy> (free, virtual and in-person)
 - z/OS Academy – Debug Explorer (2021): <https://learn.ibm.com/course/view.php?id=9899>.

- **Cheryl's:**

- ["Drink Me or Eat Me"? Is my mainframe the right size?!](#), **Jim Horne**
 - Systematic approach using SMF 70.1 (CPU by LPAR) to estimate the size of a machine you need versus the size you have.
 - Discussion afterwards brought up many points, such as those supporting the use of sub-capacity processors to get more zIIPs and lower costs.
 - Cheryl brought up egregious behavior when users specify 30-minute interval for SMF type 70 (CPU) and type 72 (service class periods).

Highlights of the Week

- **Mine:**

- Meeting Cheryl face-to-face for the first time in 3 years.
- **Andreas'** [What's New in WLM](#) session on Sunday – Excellent!
- **Brad's** and **Meral's** [WSC Hot Topics](#) session – always a Must-See at SHARE.
- Meeting *so many* people that have been helping me via emails.
- **Gary's** whistle-stop [z/OS 3.1 Preview](#) on Monday.
- [Mainframe Performance Observability Shootout](#) on Tuesday – *fantastic* job done by all the presenters – **Meral** and **Nicole** are hoping to do it again in New Orleans.
- **Mike Fitzpatrick's** [z/OS Communications Server Performance Update](#) (highly recommended by **Todd Havekost**).
- **Martin Packer's** '[So You Don't Think You're An Architect?](#)' and **Todd's** '[PSP: Collaborate and solve problems faster with unified SMF analysis](#)' sessions.
- Variety of sessions on resiliency and hacking – thank you, I won't sleep for the next week now.

- RMF CF Optimizations Enhancements.
- CPENABLE considerations
- Erase on Scratch enhancements
- Considerations for SSD drives
 - Logical erase, physical erase, UNMAP enhancements, mirroring
- Software Risk assessment & deployment considerations
- Compliance APARs
- Misc hints and tips

RMF CF Optimization

RMF CF Optimization

- Some Coupling Facility statistics are maintained in the CF, others are kept in the CFRM data sets. ADG/RMF and other monitoring tools and products use the IXLMG service to retrieve information about the CFs in the sysplex.
- IXLMG requests that gather info from a CF run synchronously (and cannot be converted to asynch by the XES heuristic algorithm), so the longer they run, and the more frequently they run, and the more systems they run on, the higher is the z/OS CPU cost.
 - In the case of ADG/RMF, it gathers information for the Monitor III data tables every MINTIME seconds (defaults to 100 - for interactive use), and)separately) every SMF interval to populate the SMF 74.4 records.
 - In addition to z/OS CPU time, these requests also consume CF CPU time and can cause serialization delays.

- **Prior to the ADG/RMF CF Optimization enhancement in z/OS 2.5:**
 - Only 1 system in the sysplex (the RMF Master) gathers the *full* set of information **for the SMF 74.4 records**. However, non-Master systems *still* gather a *subset of the* information for their 74.4 records from the CF .
 - *Gathering* of information **for the Monitor III data table** is a separate process, controlled by the CFDETAIL|NOCFDETAIL parm in ERBRMF04. The **default** is CFDETAIL, meaning the system gathers the *full* set of HW stats from the CF every MINTIME seconds *regardless of the MASTER|NOMASTER state of that system*.
 - Even worse, even though the HW stats are *gathered* from the CF on every CFDETAIL system, they are only actually *used* on the RMF Master system.

- To eliminate unnecessary data collection, RMF/ADG in z/OS 2.5 introduced a new CF Optimization option. **Note this is NOT enabled by default – you must explicitly enable it.**
- If CFOPT is specified on the S RMF or F RMF commands (or on the PARM in the RMF Started Task JCL) for z/OS 2.5 or later systems, CF hardware statistics will only be gathered on the RMF Master system.
 - Eliminates gathering of duplicate data for SMF 74.4 records on non-Master systems.
 - Eliminates gathering of duplicate data for ERBCFIG3 Mon III table on non-Master systems.
 - Pre-z/OS 2.5 systems will continue to behave as before.
- Note that specifying CFOPT on one system enables it on *all* z/OS 2.5+ systems in the sysplex, and NOCFOPT disables it on *all* systems.

- The second CF Optimization enhancement in z/OS 2.5, is the ability to specify which CF(s) each RMF instance gathers HW statistics for.
 - You might use this in a multi-site sysplex, to avoid long-running IXLMG requests to a CF in the other site (avoiding the 10 mics/km impact).
- This capability is enabled by a new CFHELP subparm on the MASTER statement in ERBRMF04 that can be used to list the CFs this system should gather CF hardware stats for.
 - For more info, see [IBM z/OS Data Gatherer User's Guide](#).
- If you use this capability, we recommend defining a 'helper' for EVERY CF. The RMF Master will collect HW stats for any CF that doesn't have an active Helper, and you don't know which site the Master will be located in at any given time.

RMF CF Optimization

SMF 74.4 sections impacted by CF Optimization

SMF 74.4 Section	RMF Master NOCFOPT	RMF Master CFOPT	Not RMF Master NOCFOPT	Not RMF Master CFOPT
Connectivity	Y	Y	N	N
Structure	Y	Y	N	N
Request	Y	Y	Y	Y (Partial ^a)
Cache	Y	Y	Y	N
SCM	Y	Y	Y	N
Async CF Duplexing	Y	Y	Y	N

^a See z/OS 2.5 SMF Manual or ERBR744 member of SYS1.MACLIB for info about which fields are valid.

- Things to watch out for:
 - You *might* consider running the Master on a ‘utility’ system (with a ‘slim’ software stack) to stop the remaining long-running IXLMG requests from blocking more important transactional workloads (and to consume cheaper MSUs if using R4HA-based SW pricing).
 - But remember that you still need to combine the SMF 74.4 records from all the systems in the sysplex to get a sysplex-wide view of CF activity.
 - *Prior* to enabling CFOPT, you might want to specify NOCFDETAIL on some systems to reduce the duplicate processing. CFOPT eliminates HW statistics gathering on non-Master systems, so *after* enabling CFOPT, specify CFDETAIL on every system to ensure ERBCFIG3 data is collected, regardless of which system is the RMF Master.

- Tips (cont) :
 - Be aware that if you specify CFHELP(xxx) on multiple systems, ALL of those systems will gather CF HW Statistics for that CF – RMF does *not* select just one Helper system as it does for the RMF Master function. You can *add* helper systems using the F RMF command, but you can't remove the helper role from a system unless you stop RMF. But this can be managed with some smart automation and rolling IPL sequences.

RMF CF Optimization

- For more information, refer to:
 - SHARE Virtual Experience 2021 session ‘[What’s New in RMF Data Gatherer](#)’ by **Peter Muench**.
 - The ‘[sysplex-wide optimized coupling facility hardware data collection \(CFOPT\)](#)’ section in the *z/OS Data Gatherer User’s Guide*.
 - ‘[Optimized RMF Coupling Facility Data Gathering](#)’ article in *Cheryl Watson’s Tuning Letter 2022 No. 3*.
- Unfortunately, there are no stats kept for the IXLMG requests, so you can’t accurately size the cost before enabling CFOPT, or the savings after you enable it.
 - However, you might see the benefit indirectly....
- **Thanks to Peter Muench in ADG Development for all his help with this.**

CPENABLE Considerations

CPENABLE Considerations

- The CPENABLE(II, hh) parm in IEAOPTxx influences how many CPs are enabled to process I/O interrupts.
 - II specifies the % of interrupts handled by Test Pending Interrupt at which WLM will *disable* a CP for handling I/O interrupts.
 - hh specifies the % of interrupts handled by Test Pending Interrupt at which WLM will *enable* another CP for handling I/O interrupts.
- The *default* value is (0,0), meaning that all VM and VH CPs are enabled for I/O interrupts.
- The *IBM-recommended* value for CPENABLE was decreased from (10,30) on z13 and previous CPCs, to (5,15) on z14 and later CPCs.
- This is described in an excellent IBM paper, '[z/OS CPENABLE Recommendations](#)' by **Kathy Walsh** and **Brad Snyder**.



CPENABLE Considerations

- The IBM recommendations are based on IBM measurements, which show that minimizing the number of enabled CPs, within reason, improves overall efficiency, delivering a 0-5% improvement in ITRR compared to the default (0,0) value.
 - Fewer interrupts for a CP should reduce opportunities for L1 cache ‘pollution’.
 - I/O Interrupt code is not taking up more L1/L2 cache than necessary.
 - The cost of bringing a waiting logical processor out of a no-work wait simply to handle an I/O interrupt has increased over time.
- But there is no free lunch – fewer CPs enabled for I/O interrupts might result in higher delays to process interrupts.
 - WLM only evaluates if more/fewer enabled CPs are needed every 20 seconds* – 20,000,000 mics is a *long* time when I/O response times are as low as 10s of mics.

- One of our readers, **Jan Tielemans**, found they were having intermittent, but noticeable, increases in response times and increased queueing. This was tracked down to long-running Db2 Commits (> 1 second). The commits were experiencing occasional long response times due to delays in processing the I/O interrupts.
- They run a long-distance multi-site sysplex, with workloads and CFs distributed across the two sites. After *much* dump analysis, they determined that the interrupts were being delayed because:
 - Very long-running IXLMG requests were tying up the CPs that were enabled for I/O interrupts for hundreds or thousands of mics, causing long queues of I/O interrupts waiting to be processed.
 - With only 2 CPs enabled for interrupts most of the time, there is an increased chance that one or both of the enabled CPs will be busy processing long-running IXLMG commands.

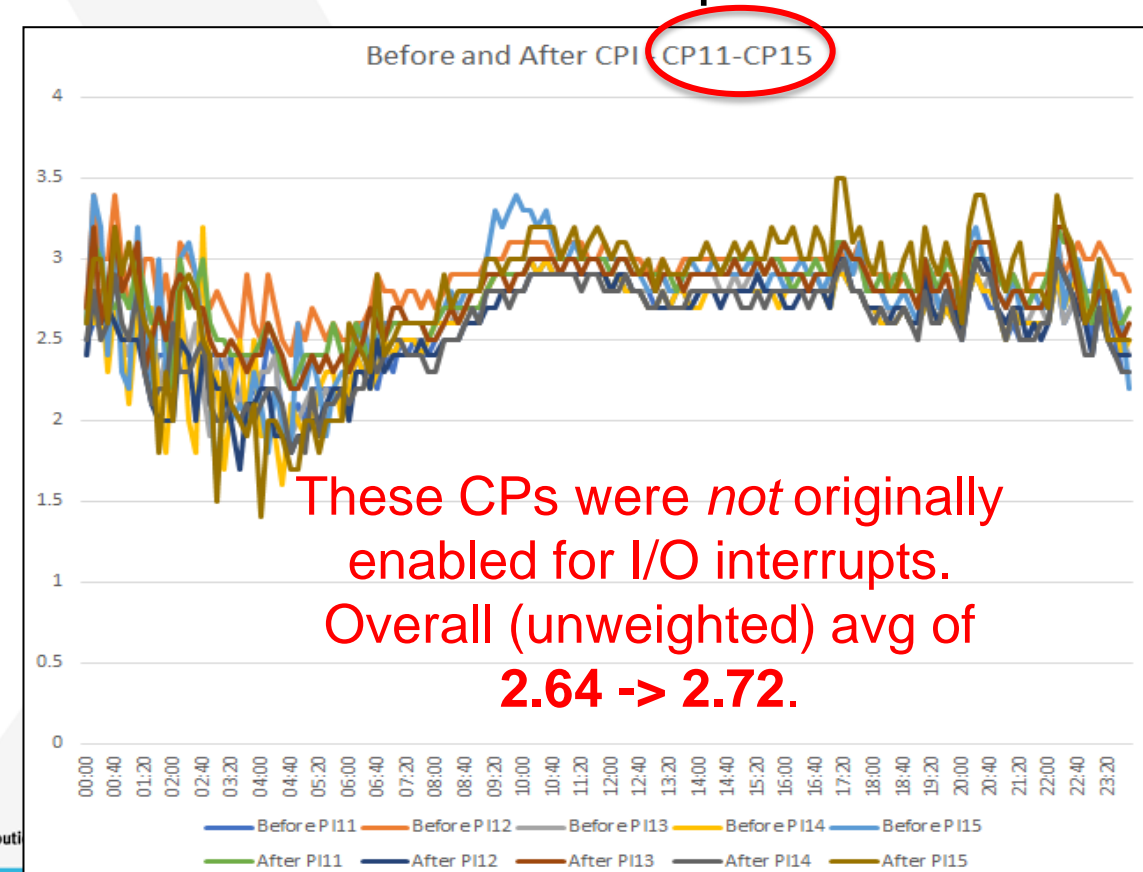

CPENABLE Considerations

- Jan did some interesting SMF analysis.
 - With CPENABLE set to the recommended 5,15, the average delay during the prime shift was around 0.3-0.5 ms.:
- With CPENABLE set to 0,0:
 - You can see the dramatic reduction in interrupt delay time (the highlighted drop is where the CPENABLE value was changed).
 - No more Db2 commit times > 1 sec.




CPENABLE Considerations

- The aforementioned IBM paper clearly states that reducing the CPENABLE values could reduce the Interrupt delay time, but the price for that could be a worse ITRR. What was the impact on CPI of changing CPENABLE in Jan's shop?



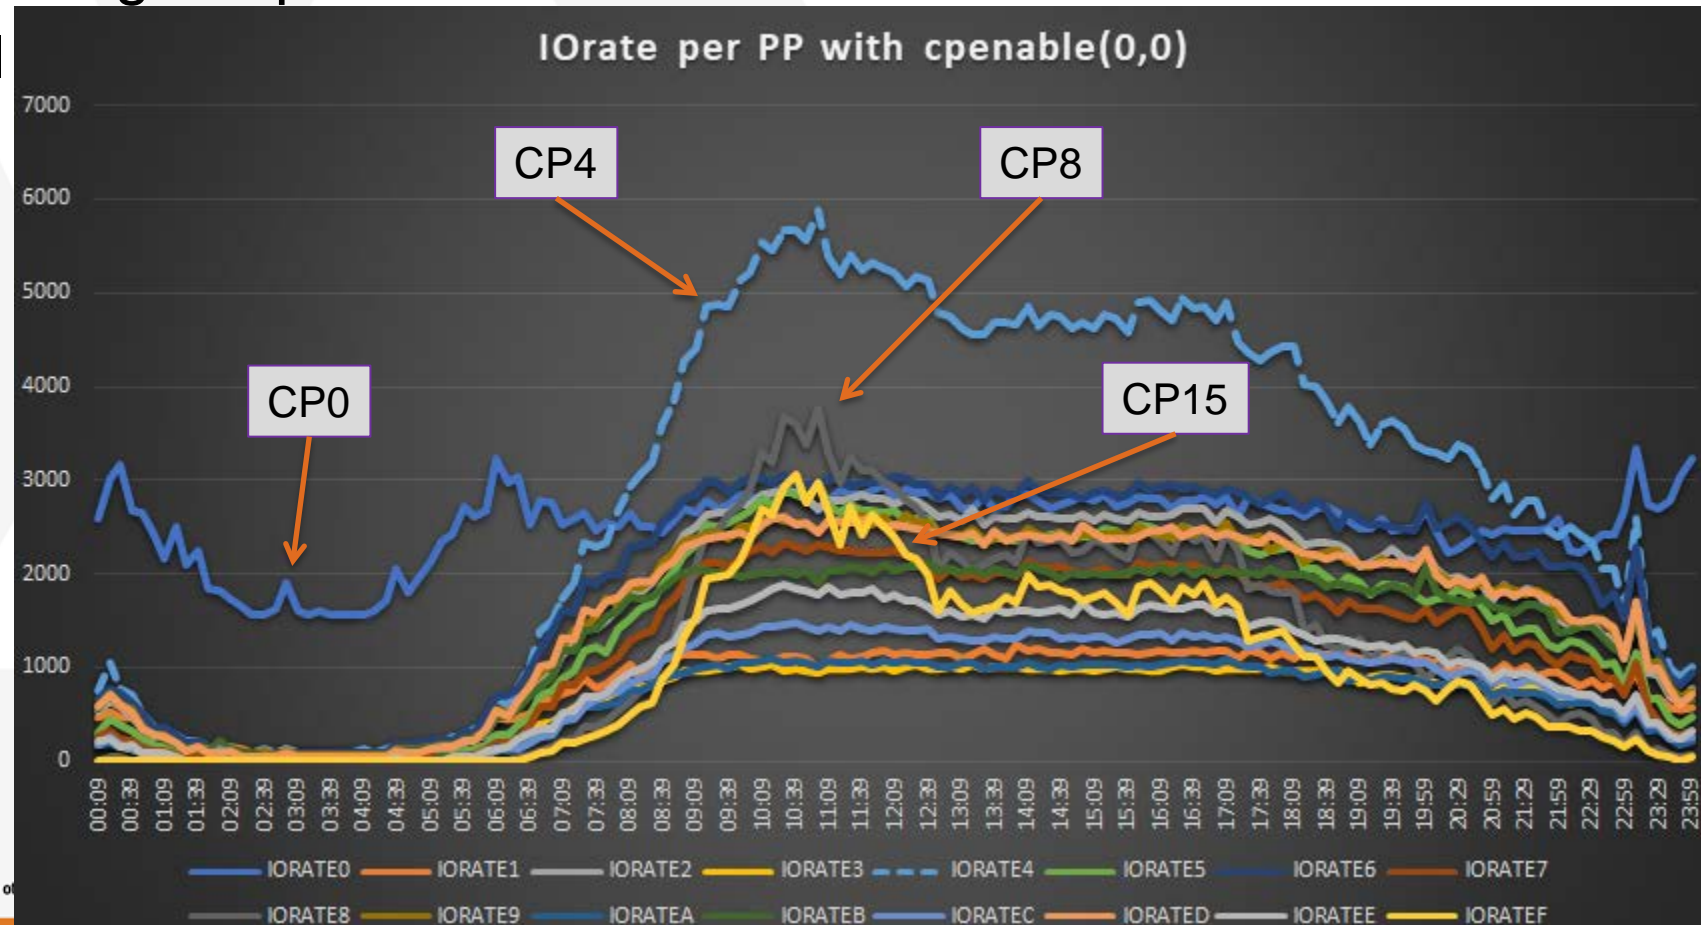
CPENABLE Considerations

- There was also an interesting side effect of enabling additional CPs for interrupt handling.
 - With (5,15), the bulk of interrupts were handled by CP0, with it consistently handling more interrupts than CP1.



CPENABLE Considerations

- With all CPs enabled for I/O interrupts (COENABLE(0,0)), the bulk of interrupts were handled by CP4 during the prime shift.
 - CP0 was 'normal' until the rate increased.
 - Most of the other CPs behaved consistently except for CP15, the only VM CP.



CPENABLE Considerations

- Bottom line is that CPENABLE works as expected.
- However, in this case, reverting to (0,0) addressed the issue of high interrupt delay times, and without any discernible impact on CPI or MSU consumption.
- If you experience intermittent long I/O response times, especially if you have the possibility of very-long-running IXLMG requests, consider decreasing the CPENABLE values to see if that helps.
 - Enabling RMF CF Optimization potentially reduces the number of long-running IXLMG requests, so try enabling that (if possible) before changing CPENABLE.
 - Both CPENABLE and RMF CFOPT can be changed dynamically.
 - Don't forget to keep an eye on CPI to make sure any changes are acceptable.
- Thanks to **Jan Tielemans** and **Dale Riedy** for their help with this.

Erase-on-Scratch (EoS) and Related Topics

- Erase-on-Scratch was introduced *many* years ago, to address the potential exposure of residual data being left on a disk after a data set is scratched.
 - Tracks overwritten with zeros as part of Delete processing. Seemed like a good idea. But the performance was *dreadful* – I don't know anyone that turned it on in production at the time.
- Over subsequent years, changes were made to reduce the exposure (e.g. by writing an EOF marker at the start of newly-allocated data sets) and to improve the performance of overwriting the tracks where the data set previously existed:
 - In z/OS 2.1, the code that overwrites the tracks was changed to write 255 tracks per EoS Channel Program (12,240 in z/OS 2.2), rather than 1 track/CP previously. This reduced both the number of EXCPs and the elapsed time compared to z/OS 1.13.
 - By default, these enhancements only apply to data sets on non-mirrored volumes. To have them apply to mirrored volumes, you **must** specify EOSV2=YES in DEVSUPxx.

- Despite those enhancements, many customers *still* didn't enable Erase on Scratch.
- Then auditors heard about the exposure and how to address it (see [STIG V-259](#) and [z/OS RACF STIG](#)). Now we are seeing “Is EoS enabled?” showing up on external auditors' checklists. (See [this session](#) for more about STIGs.)
- The good news is that, as long as you don't have jobs that do nothing but delete large numbers of very large data sets, the impact of enabling EoS for most jobs should be relatively small compared to the day-to-day elapsed time variability of those jobs.
 - Just make sure that you specify EOSV2=YES if you mirror your disks.
 - If you use non-IBM disk subsystems, ask your vendor for their recommended values for the LRECCW_PER_EOSCP parm in DEVSUPxx.

- You can enable EoS for:
 1. ALL data sets.
 2. Data sets with a RACF profile SECLEVEL above a given value.
 3. Data sets that specify ERASE in their RACF profile.
- You *could* use Option 3 to minimize the impact of using EoS by limiting the number of 'EoSed' data sets. However you are exposed if someone copies an 'ERASE' data set to a data set whose profile does *not* specify ERASE – when *that* data set is deleted, the residual data is not overwritten.
- Option 1 is the most secure of the 3 options. It is also the only one that protects temporary data sets*. But it also has the greatest performance impact because the residual data for *every* data set needs to be 'addressed'.

* More on temp data sets later


EoS Considerations for SSDs

- On HDDs, Erase-on-Scratch would physically write zeros to the disk when a data set is deleted or has a PARTREL issued for it.
 - This protects you from a user on your system being able to read residual data.
 - It *also* protects you when the HDD goes out of your control - for example, when you decommission the disk subsystem.
- SSDs work differently. All writes must be to a blank page in the SSD, with the page that is being updated being marked as Invalid.
 - Erase-on-Scratch with SSDs would protect you from a user on your system being able to read residual data.
 - However, with the right equipment, someone might be able to read residual data in the 'Invalid' pages that have not yet been recycled if they get their hands on the SSD. This should be borne in mind when decommissioning an SSD subsystem.

EoS Considerations for SSDs

- However, there is another concern for SSDs. SSDs have a finite number of “program/erase” (P/E) cycles before they need to be replaced. So if you enable Erase-on-Scratch for a large number of data sets, that could result in a much-higher-than-normal number of writes to the SSD. This is *not* a good thing.
- To address this, APAR [OA61492](#) (for z/OS 2.4 & 2.5 with DS8900) delivered a new UNMAP function. Rather than overwriting the data, UNMAP tells the disk subsystem to mark those pages as Invalid – Invalid SSD pages cannot be accessed from a program on any IBM operating system.
 - This protects the data and also addresses the finite P/E cycles issue. ✓
 - As an added benefit it also delivers performance enhancements because the work to mark the pages as invalid is done by the disk subsystem and makes the performance impact of EoS nearly invisible. ✓

Elapsed Time (Secs) for Delete of 63K Track Data Set



EoS Considerations for SSDs

- But (there's always a 'but'... 😞)
- UNMAP is **currently not supported** for devices in a copy services relationship.
 - Erase-on-Scratch processing for data sets on those SSD devices will revert to the traditional overwrite-with-zeros model, so residual data *is* protected.
 - But it eliminates the performance benefit of UNMAP for those devices.
 - If you are using EoS with HDDs today, moving to SSDs that are in a Copy Services relationship should result in similar elapsed times.
 - It exposes you to the issue with the finite number of P/E cycles.
- If this is an issue that would impact you, watch out for an upcoming 'Idea' to add mirroring support for UNMAP and vote in favor.
- See Section 3.7.3 of [this Redbook](#) for info about protecting any potential residual data when decommissioning SSD devices.

- Misc other tidbits:
 - If you have the UNMAP SAF FACILITY class profile defined and do *not* have SETROPTS ERASE(ALL) specified, temporary data sets on SSD devices *will* use UNMAP if possible. However, if UNMAP can't be used (if the data set resides on a mirrored volume, for example), the data set will *not* be overwritten. This allows you to measure the benefit of UNMAP without the risk of impacting data set scratch times because of overwriting the data set's tracks.
 - Note that this is a change to the default behavior.
 - However, why would you NOT want to erase residual data if there is no performance impact from doing so?
 - The new behavior for temporary data sets is mentioned in the last paragraph in the [Unmapping DASD data](#) section of *z/OS DFSMS Using Data Sets* manual.

- Misc other tidbits:
 - The Volume Metrics Section of the SMF 42.5 records contain 6 new fields about the use of Erase on Scratch and the DADSM UNMAP function. Additionally, the type 42 subtype 27 (VTOC DSCB Update) section has been updated to include EoS/UNMAP information. Definitely worth looking into.
 - If forced to enable EoS, but performance impact on a DELETE/CREATE job is unacceptable, consider RENAME/CREATE, then DELETE the renamed data set asynchronously in a separate job.
- An alternative approach to achieve a similar net effect might be to use data set encryption.
 - Residual data is not overwritten, but it *is* encrypted/meaningless.
 - However, data set encryption is data set-level vs. ERASE ALL, which covers all data sets.

EoS Considerations for SSDs

- If you *can* use UNMAP, we don't see any reason *not* to enable ERASE ALL.
- If you *can't* use UNMAP, we recommend preparing yourself for an EoS discussion by running some EoS measurements in your environment to determine the precise impact of your HW and your distance (EoS cost scales with data set size). Then analyze SMF records to identify the number and distribution of deletes of large data sets. Apply results of measurements to determine elapsed time impact.
- The decision of whether or not to use EoS should factor in the projected performance impact, but ultimately, this should be a security decision, not a technical one.

- **What would you think about changing the EOSV2 to default to YES?**
- For more information about Erase on Scratch, refer to:
 - [z/OS DFSMS Using Data Sets](#) manual.
 - [z/OS RACF System Programmer's Guide](#) manual.
 - '[Erase on Scratch](#)' article in *Cheryl Watson's Tuning Letter 2015 No. 1*.
 - '[Proposed Erase-on-Scratch Enhancement](#)' article in *Cheryl Watson's Tuning Letter 2019 No. 3*.
 - US Defense Information Security Agency [IBM z/OS RACF Security Technical Implementation Guide](#).
 - '[Erase-on-Scratch In the Modern World](#)' NASPA article (from 2000!) by **Mark Nelson** and **Wayne Rhoten**.
- Thanks to **Carrie da Silva** in DFSMS Development for all her help and enthusiasm.

Software Risk Assessment Deployment Considerations

- Software Risk Assessment is a new focus area for auditors and resiliency managers.
- It is built on information in your SMP/E CSIs to check your software inventory for:
 - Unapplied Security/Integrity service
 - Unapplied HIPER service
 - Products that are approaching (or exceeded!) their end of service date
 - Products that are approaching their end of life date
 - Assessing general service level of software products
- See '[Game Changer: Software Risk Assessment](#)' SHARE in Atlanta 2023 Session by **Todd Gagle** and **Dawn Damore** and Todd's '[Software Risk Assessment](#)' article in Cheryl Watson's Tuning Letter 2022 No. 3 for more information.

- The Software Risk Assessment is valuable and worthwhile because it puts a more formal process in place for reducing the risk of un-addressed exposures on your system.
- However, it is challenged by the fact that SMP/E's role ends when it has updated your target libraries.
 - SMP/E has no knowledge of which systems are actually using the target libraries associated with any particular CSI.
 - The running system has no knowledge about which CSI was used to create the libraries it is using.
 - It is left to the installation to have well documented, enforced, naming conventions, policies, and cloning procedures to allow an SME to somehow associate SMP/E target zones and running systems.

- While most auditors are happy to accept the accuracy of the information in SMP/E, convincing them that a home-grown, partially manual, partially documented, process guarantees they can tie the software used by a system back to the SMP/E target zone it was built from might be a challenge.
- We recently worked with IBM Development and IBM sysprogs to evaluate the z/OSMF Deployment function against:
 - Their existing home-grown procedures, created over 30+ years.
 - Whether the ‘audit trail’ that the Deployment function provides would be more acceptable to a Software Risk Assessor.
- The net result was that it covers most of the functions provided by the current processes, and the documentation and audit trail it provides is an improvement over the current processes.

- The largest remaining concern was still that you don't have a 'hard' pointer from the running system back to the SMP/E CSI and vice versa.
- However, the z/OS 3.1 preview states:
 - “z/OS 3.1 plans to support generating a Universally Unique Identifier (UUID) during the installation and deployment of z/OS, correlating the UUID with the z/OSMF Software Management software instance. This represents that copy of z/OS, storing the UUID value in a well-known location on z/OS, and retrieving the UUID from a running z/OS system to directly identify the proper software instance.”
- If Software Risk Assessment is a hot button topic in *your* shop, evaluate the z/OSMF Deploy Software Instance function to help bridge that gap between SMP/E Control and your running systems.

Software Risk Assessment

- For more information, see:
 - SHARE in Atlanta Session '[A Tour of z/OSMF Software Management](#)' by **Kurt Quackenbush**.
 - SHARE in Atlanta Session '[Game Changer: Software Risk Assessment](#)' by **Todd Gagle** and **Dawn Damore**.
 - Cheryl Watson's Tuning Letter article '[Software Risk Assessment Considerations for Software Deployment](#)'.

‘Compliance’ SMF Records

- Has anyone noticed a number of New Function APARs over the last year, delivering new SMF Type 1154 (Compliance) records?
- This section focuses on what those records are, which products produce them, what you have to do to get them.
 - How they are used is multiple sessions all on its own – our focus here is on the supporting infrastructure.
- The infrastructure for creating these records is provided by APARs [OA61443](#) (CEA), [OA61444](#) (SMF/ENF), and [PH37308](#) (z/OSMF) for z/OS 2.4 and 2.5.
 - Requires z15 or later CPC.
- The minimum release for producing the different subtypes depends on the specific exploiting products.

- In response to the relentless flow of companies announcing that they have experienced a 'data breach', governments and industry bodies around the world are busy rolling out new regulations (for example, PCI DSS and NIST SP800-53), and harsh penalties for companies that do not meet those regulations.
 - In many breaches, the hacker is exploiting a known exposure, whose fix was available for *months* but never applied.
- The intent of these regulations is to protect businesses and individuals by making it more expensive for companies to ignore security issues than to address them.
- However, the challenge for companies is how do you ensure, and prove, that your systems and subsystems adhere to the myriad regulations now *and* at any time in the past?

- The 'Compliance' APARs deliver new functions in products to gather information related to selected regulations and standards, and save that information to SMF records.
 - Having the information should help companies identify and address breaches in a timely manner, but having a clear audit trail of the status of the system on an ongoing basis is *very* important – in many breaches, there is not a clear picture of when the system was first penetrated.
- So, which products/functions currently support this Compliance data gathering?

z/OS Compliance

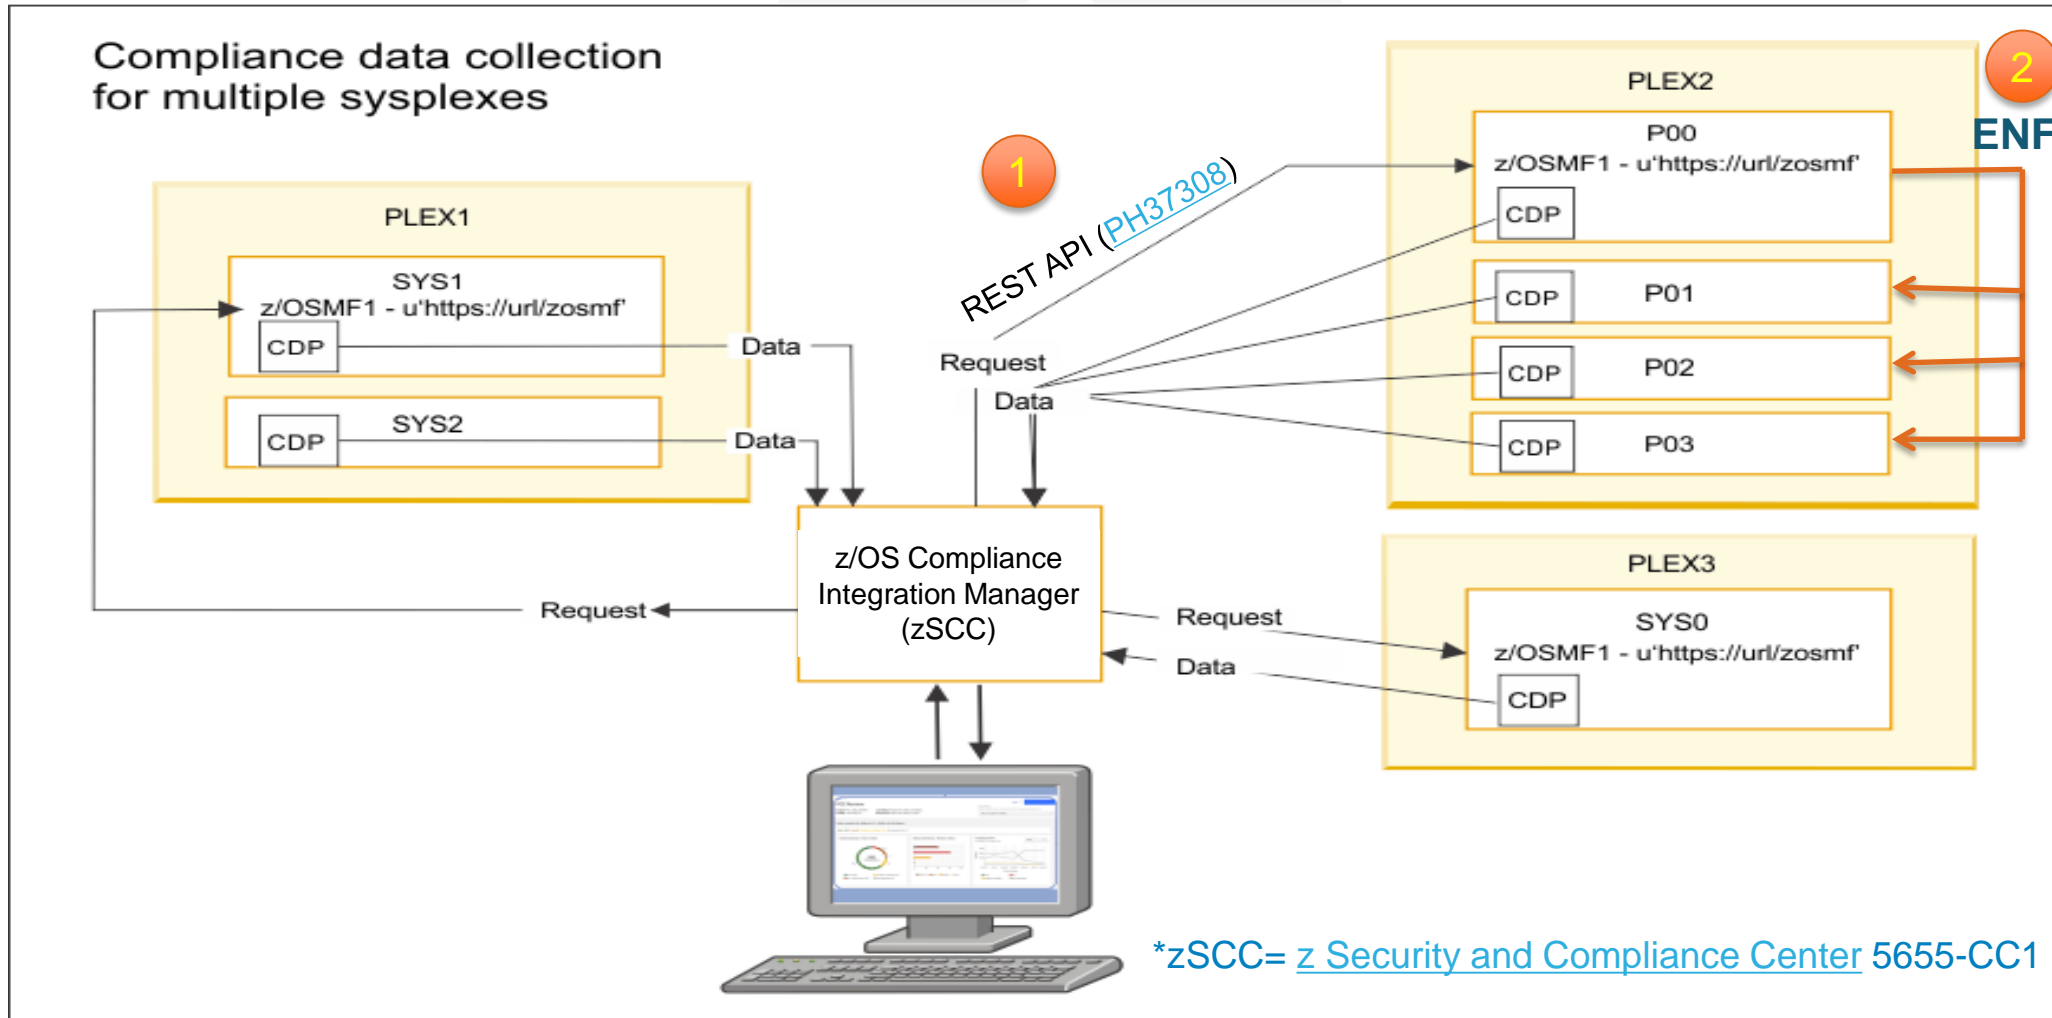
SMF 1154-producing products as of 3/8/23

1	TCP/IP stack	✓	Here	78	SSHD		
2	FTP daemon	✓	Here	79	InetD		
3	TN3270	✓	Here	80	CICS TS for z/OS	✓	Here
4	CSSMTP	✓	Here	81	Db2 for z/OS	✓	Here
49	ICSF			82	IBM MQ region		
50	Consoles			83	RACF	✓	Here
51	DFSMSdfp			85	IMS control region and IMS OTMA		
52	DFSMSrmm			86	IMS operation manager		
53	DFSMShsm			87	IMS Connect		
54	DFSMSdss			96	SMF global reporting options		
77	z/OS UNIX System Services			97	SMF subsystem reporting options		
				128	Processor activity	✓	Here

- You can get the *current* list of functions that create SMF 1154 compliance records from the [SMF Manual](#).
- The IBM.Function.Compliance.DataCollection [FIXCAT](#) should contain a list of APARs that deliver or enhance the infrastructure for this capability, or deliver exploitation in supporting products.
- The '[z/OS Compliance Data Collection](#)' manual contains the mapping of *some* of the subtypes. But for many others, it points you at the [zSCC manual](#), even though that manual does not contain the mapping for those records.
 - We are currently working on finding this information.
- How do you turn on collection of these records, and how do you report on them?

Compliance Data Collection BIG Picture

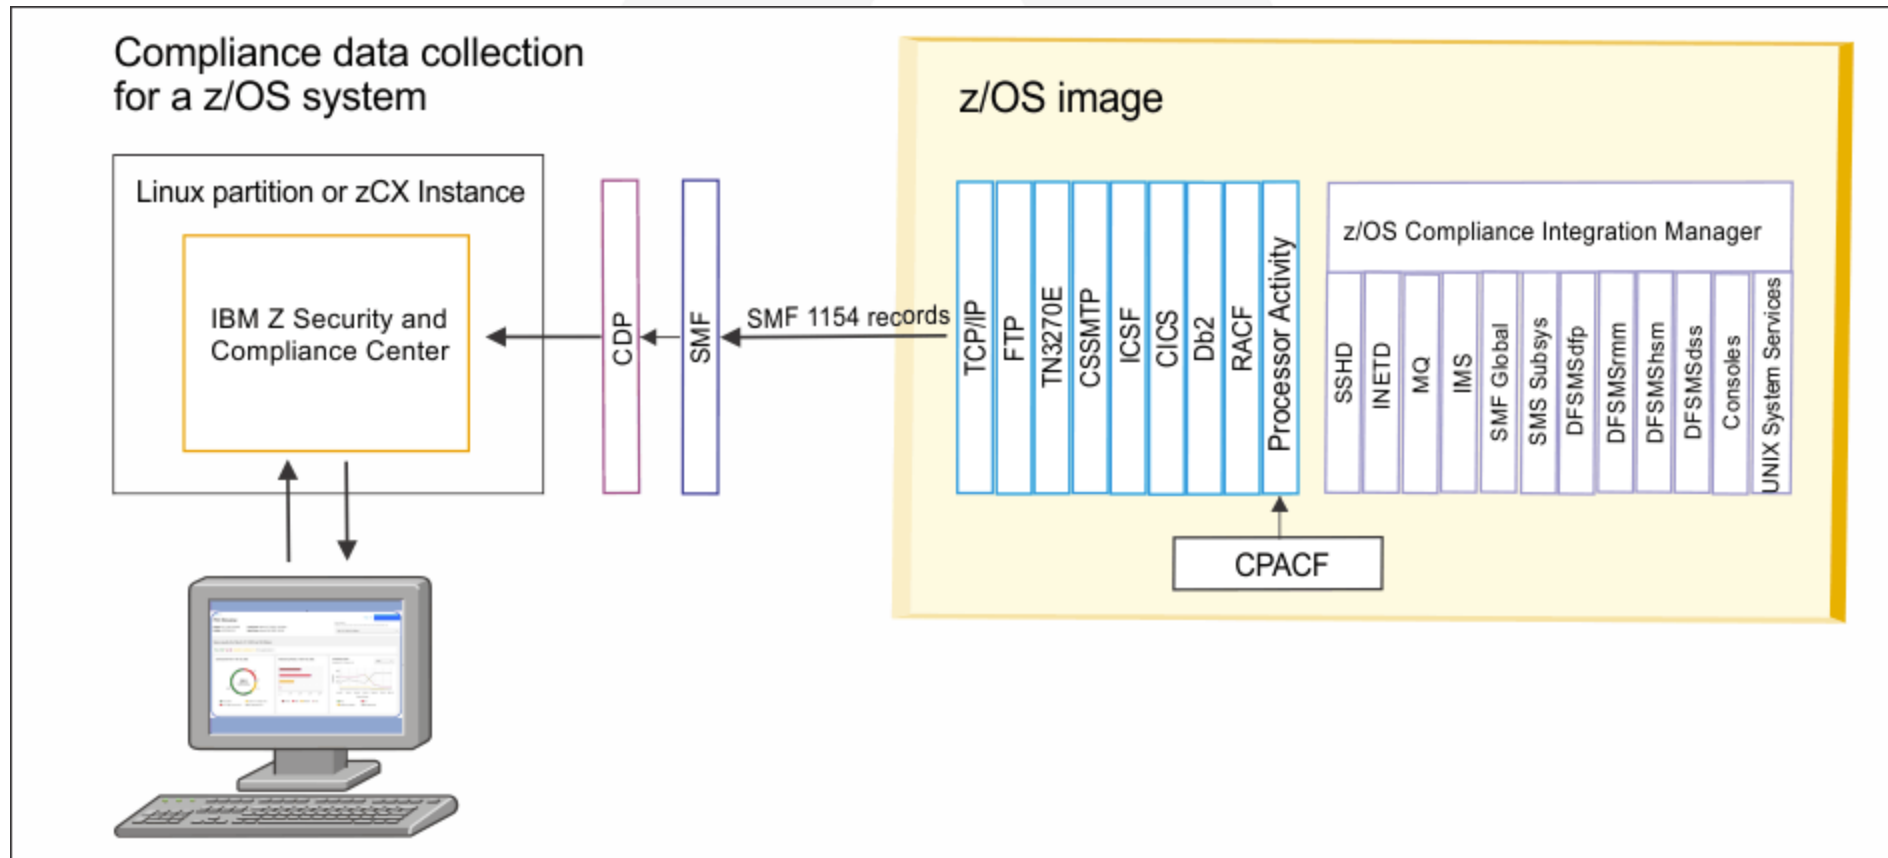
Copyright IBM Corporation



*zSCC= [z Security and Compliance Center 5655-CC1](#)

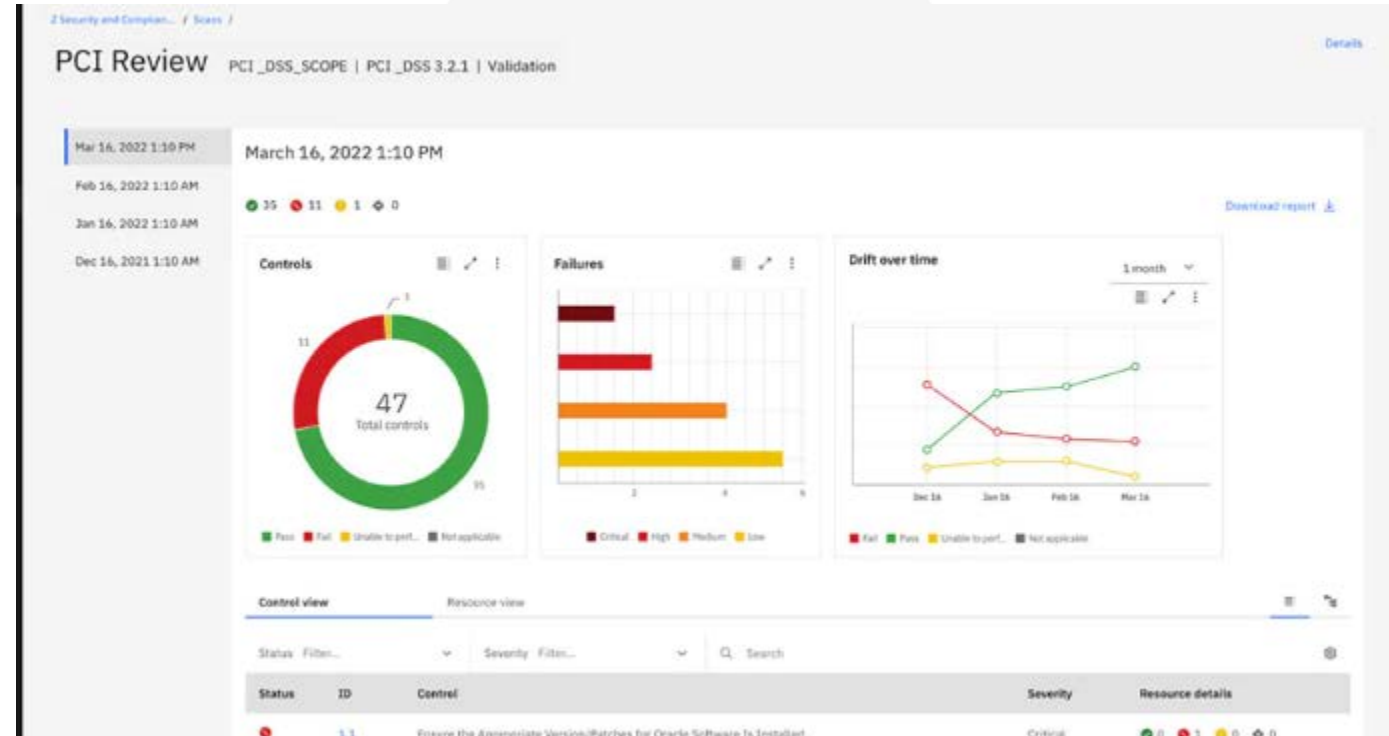
SMF 1154 Record Creation – a Little More Detail

See <https://www.ibm.com/docs/en/zscsc/1.1.0?topic=center-architecture-technology>



Copyright IBM Corporation

- The only IBM product that drives the creation of the type 1154 records is the z/OS Compliance Integration Manager component of the z Security and Compliance Center.
 - However, the REST API that it drives is a published interface, so you *could* potentially create your own program to start creating type 1154 records *now*, so you would have a history that could then be analyzed when you have an analysis tool in the future.
- IBM Common Data Provider is used to get the SMF records back to zSCC.
 - If you don't already have CDP, zSCC includes a Limited Use License for CDP.
- zSCC currently requires OpenShift on z (either Linux on Z, or under zCX).
 - If OpenShift is not part of your strategy, but you *would* like zSCC, speak to IBM.



- For more information about the IBM z Security and Compliance Center, see:
 - SHARE in Atlanta Session 84053, '[BYOD Lab: IBM Z Security and Compliance Center](#)', by **Eva Yan** and **David Rossi**.
 - IBM zSCC Home Page: <https://www.ibm.com/products/z-security-and-compliance-center>.
 - IBM zSCC Documentation: [IBM Z Security and Compliance Center 1.1.0 - IBM Documentation](#).
 - List of announced SMF type 1154 subtypes: [z/OS SMF System Management Facilities](#).

Miscellaneous Tidbits

RMF-related Changes

- If you are working on optimizing your HiperDispatch topology, it is *very* helpful to see where logical CPs are assigned. This information is available in the SMF 99.14 records, which is great. However, many sites suppress all SMF type 99 records, meaning that this information is not available to them.
- To make this information more widely available, APAR [OA62064](#) enhances ADG/RMF to include this information FOR ALL LPARS ON THAT CPC in the type 70 records if z/OS is running on a z16.
- This additional information in the type 70 records is used to create the z16 Topology Window in zPCR.

No.	Name	Type	Drawer 1								Drawer 2								
			DCM 1		DCM 2		DCM 3		DCM 4		DCM 1		DCM 2		DCM 3		DCM 4		
			Chip 1	Chip 2	Chip 1	Chip 2	Chip 1	Chip 2	Chip 1	Chip 2	Chip 1	Chip 2	Chip 1	Chip 2	Chip 1	Chip 2	Chip 1	Chip 2	
1	SSA	GP																	
	SSA	zIIP	(2)L									(4)L	(3)L	(2)H (1)M	(8)H	(8)H	(8)H	(8)H	(8)H
2	R76	GP										(4)L	(3)L	(2)H (1)M	(8)H	(8)H	(8)H	(8)H	(8)H
	R76	zIIP	(2)L											(1)H (1)M (1)L	(2)L				
3	R71	GP	(2)H	(8)H	(8)H	(8)H	(6)L	(4)H (2)M (1)L	(8)L	(8)L				(1)L	(1)L				
	R71	zIIP	(1)H (1)M (1)L											(1)L	(1)L				
4	R77	GP										(4)L	(4)H (1)M	(3)L	(8)L	(4)L			
	R77	zIIP												(1)H (1)M					
5	S5D	GP	(2)L			(2)L	(6)L	(7)L	(8)L	(4)L									
	S5D	zIIP	(3)L											(3)L	(2)L				
6	S5H	GP												(3)L	(2)L				
	S5H	zIIP	(3)L											(2)L	(1)L				
7	TA4	GP												(2)L	(1)L				
8	S5B	GP																	
	S5B	zIIP	(3)L											(2)L	(1)L				
9	X7CFA89	ICF												(2)L	(1)L				
10	X5CFA89	ICF																	

Key to Above Table
 -: Not Polarized H: Vertical High L: Vertical Low M: Vertical Medium

Copyright IBM Corp.

RMF-related Changes

- Also related to RMF is announcement [221-338](#), stating that the z/OS Workload Interaction Correlator is entitled (FREE) to licensed RMF or ADG customers.
 - Previously this was a chargeable z/OS feature code (see [OA62268](#)).
- And don't forget recent best practice '[Always Collect Correlator SMF Records](#)'.
 - This relates to the HFTSINTVL parm in SMFPRMxx – default is NOHFTSINTVL.
- If you are interested in zWIC or zWIN, see **Dan Rosa's** '[Diagnosing Transient Performance Issues by Analyzing SMF 98 and 100 Records in Zowe](#)' at 4pm in Hanover D.
- There is also a [zTrial](#) available for zWIN. Note that it does NOT have the latest enhancements – speak to Dan if you would like more information.

RRS/Logger Tip

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

JES2 Spool Compression

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

- Thanks to **Tom Wasik's** [*I Didn't Know You Could Do That With JES2!*](#)
 - Don't use MASDEF HOLD and DORMANCY! Instead use MASDEF CYCLEMGT=AUTO to allow JES2 to automatically adjust them.
 - **Significant** performance improvement for SYSOUT work selection by using a binary sort for posting and selecting. Specify OUTDEF LDEV_OPT, SAPI_OPT, and OUTDEF WS_OPT.
 - To improve performance, be careful of the DEBUG statement. QVERIFY, VERSIONS, CKPT use lots of CPU.

(Non)Reader Beware

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

This is important

(Non)Reader Beware

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

(Non)Reader Beware

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

Accessing IBM FTP sites

- IBM Support Transfer ID
 - I recently tried to look at the IBM z/OS New Function APARs website, but found that you now need an 'IBM Support Transfer ID' to access that and other IBM FTP sites.
 - I was dreading IBM bureaucracy, but with help from **Marna Walle**, I was up and running in a couple of minutes!
 - Go to: <https://www.secure.ecurep.ibm.com/transferids/>. You will need to sign in using your traditional IBM ID and password, but once you do that, you are immediately presented with a Transfer ID and Password.
 - Make sure to save them, they are not very memorable, and not recoverable.

One for the road

- If someone initialized a production 3390-54, how would you:
 - Figure out what data sets were on the volume?
 - Find the most recent backup of each one?
 - Handle multi-volume data sets with one extent on that volume?
 - Recover if the volume contained a user catalog?
 - How about if it is the user catalog that holds your archived SMF files?
 - Recover if that volume contained your HSM (or equivalent) control data sets?
 - Recover if that volume contains your security database?
 - Recover if the volume contained a spool data set? How would you know which outputs were in the missing data set?
 - Bring all those data sets up to the point in time before the VTOC was trashed?
 - Do all this before the CEO calls to see what is going on?

One for the road

- Who thinks 'ransomware' doesn't apply to z/OS?
 - After all, encryption is at the data set level, and not all data sets support encryption.
 - And consider how long it takes to fully implement data set encryption.
- If you *still* suffer from a lack of sleepless nights, consider attending the no-charge [IBM Resiliency Conference](#) in Munich from June 20th to 23rd.

**End
Thank You!!!**

Your feedback is important!

Submit a session evaluation for each session you attend:

www.share.org/evaluation

