

Session z104139

The Watson & Walker zRoadshow

—

Cheryl Watson, Frank Kyne

cheryl@watsonwalker.com

Watson & Walker



TechU

2018 IBM Systems Technical University

October 11, 2018

Hollywood, FL

IBM®

Welcome!


- Cheryl Watson, President of Watson & Walker Inc. since 1986, working on IBM mainframes since 1965
- Frank Kyne, Editor and Technical Consultant with Watson & Walker since 2014.
- We publish Cheryl Watson's Tuning Letter (since 1991), teach classes, provide consulting, and have three software products: BoxScore, BoxScore II, GoalTender.
- Our latest passion is our SCRTPro Service Offering, where we help clients get the maximum value from their z/OS systems.
- We've also expanded our team to include Brenda White, Mario Bezzi, and Alan Murphy!
- We are here to talk a little bit about a ~~lot~~ few things we are think important.



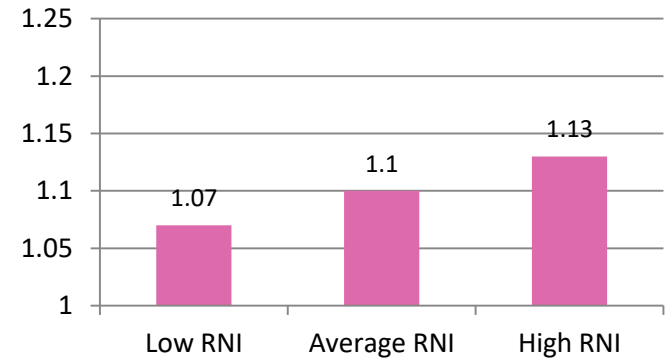
z14 Tips and Early Experiences

Generation to generation relative performances

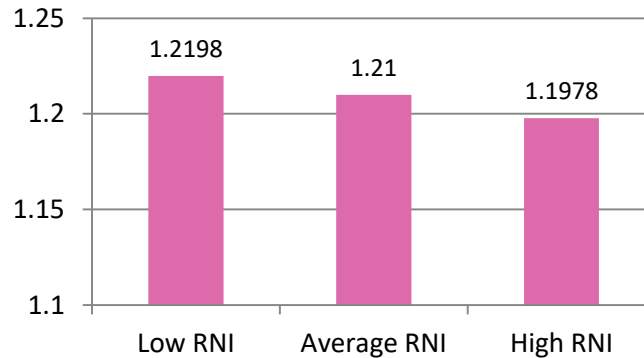
zEC12 vs z13



z13 vs z14



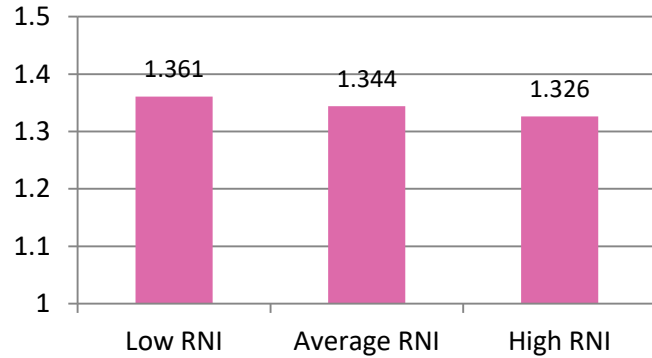
zEC12 vs z14




z14 Tips and Early Experiences

Generation to generation relative performances


zBC12 vs z13s



z13s vs z14 ZR1



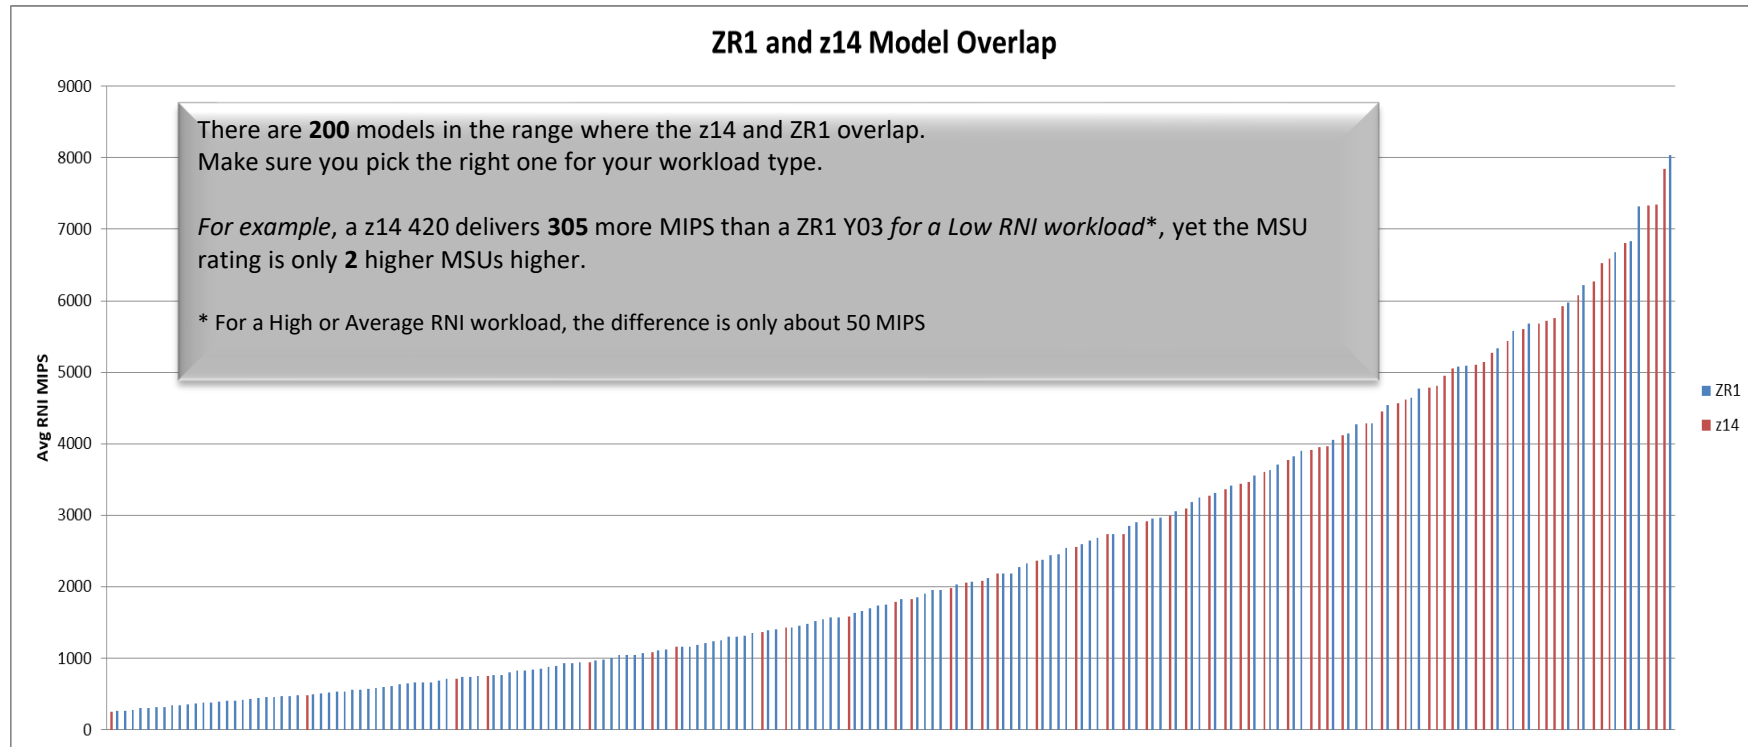
zBC12 vs z14 ZR1



* Based on Z01 numbers

IBM ZR1

You have all heard about IBM's new ZR1. We just wanted to point out a few interesting things:



z14 Tips

When you run zPCR, make sure that you always download the latest level of CP3KEXTR and include the type 113 SMF records.

The following JCL illustrates how to create the EDFs for multiple systems in a single run:

```
//EXTR      EXEC  PGM=LOADER
//SYSLIN    DD   DISP=SHR,DSN=YOUR.CPSTOOLS.JCL(ZOBJEXTR)
//SYSOUT    DD   DUMMY,SYSOUT=*
//SYSLOUT   DD   DUMMY,SYSOUT=*
//SMFIN     DD   DISP=SHR,DSN=YOUR.SMF.INPUT.FILE
//EDF001    DD   DISP=SHR,DSN=YOUR.EDFI.OUTPUT.FB80(SYSA)
//PRINT001 DD   SYSOUT=*
//EDF002    DD   DISP=SHR,DSN=YOUR.EDFI.OUTPUT.FB80(SYSB)
//PRINT002 DD   SYSOUT=*
//SYSIN001  DD   *
ENT='CUSTOMER DATA'  YOUR COMPANY NAME (REQUIRED)
SYSID=SYSA            SMF SYSID (REQUIRED)
DURATION=01:00       DURATION MUST BE A MULTIPLE OF RMF INTERVAL
SORT=YES             IF SMF RECORDS ARE NOT IN SORTED ORDER
/*
//SYSIN002  DD   *
ENT='CUSTOMER DATA'  YOUR COMPANY NAME (REQUIRED)
SYSID=SYSB            SMF SYSID (REQUIRED)
DURATION=01:00       DURATION MUST BE A MULTIPLE OF RMF INTERVAL
SORT=YES             IF SMF RECORDS ARE NOT IN SORTED ORDER
```

Did I mention that
you should capture
AND KEEP your SMF
113 records?

It's All About the ~~money~~ Cache

One of our customers, **NASCO**, recently upgraded from a z13 709 (12,371 Average RNI MIPS) to a z14 523 (12,584 Average RNI MIPS).

Following the upgrade and configuration optimization, their peak R4HA MSUs dropped by 22%!

– **NOTE: NOT EVERYONE WILL ACHIEVE THE SAME RESULTS.**

It is well known that additional *capacity* reduces the MSUs to do a given amount of work. (See [USAA's experience](#) in [Tuning Letter 2015 No. 4](#) and **Todd Havekost's** award-winning SHARE presentations on "[Achieving CPU \(& MLC\) Savings Through Optimizing Processor Cache](#)" for more info.) This example shows what might be achieved by having additional *CPs* and *more cache, more memory* (640GB to 2112GB), and a workload that likes that generation (NASCO had a High RNI Workload), but the 'same' capacity.


We are hoping to document their experiences in an upcoming Tuning Letter issue.

Encryption

- You probably didn't notice, but IBM also had a low-key announcement for a new set of capabilities called Pervasive Encryption.
- Why you should care - because prison or HUGE fines are Not Nice Things.
 - AND because your Executives probably care.
- You have a few things to consider:
 - Performance/overhead.
 - » IBM measurements show average overhead of 11.6% on z13 vs 2.6% on z14. Naturally, YMMV.
 - » SMF 42.6 records have been enhanced ([OA52132](#), [OA52133](#), [OA54663](#)) to allow zBNA (1.8.1) to provide very accurate data set-level predictions (called 'projection support').
 - » Data is decrypted/encrypted when it moves in or out of z/OS. So I/O-intensive workloads experience higher-than-average overhead. Seems like this would be an ideal application of the benefits of large memory. However, IBM have not performed any measurements yet to determine whether this is the case, or to what extent.

Encryption

- Few things to consider
 - Key management
 - This is a critical topic, especially if you would like to be able to access your data x years in the future.
 - For more information, see **Andy Coulson's z/OS Dataset Encryption - Why and How, Session z100003.**
 - Start getting ready NOW.
 - PDSs are not and will not be supported by encryption.
 - Non-Extended Format sequential and VSAM data sets are not supported.
 - Don't wait until your CIO asks you to "flip the Pervasive Encryption switch". At a minimum, stop the challenge from getting even worse by switching new allocations to use PDSEs and EF sequential and VSAM data sets NOW (also required for zEDC).
 - Enable CPU MF Crypto counters on HMC and F HIS command
 - For an up-to-date list of unsupported data set types, see ???
- For a guide for those of us that *don't* have super computer brains, see draft Redbook *Getting Started with z/OS Data Set Encryption*, [SG24-8410](#) (last updated on May 4).
- If you would like a very simple, step-by-step guide to getting started with Encryption, refer to [this website](#) by **Eysha Powers**.



Sub-capacity software pricing

- At a recent conference, one of the presenters said that his company spends more on software upgrade fees than they do on the hardware when they upgrade their CPCs.
- Software licenses based on full capacity make no sense.
 - Nearly all vendors that we are aware of will agree to charging for software on a sub-capacity basis (like z/OS) either voluntarily, or if pressed.
 - But the time to negotiate this is NOT on the last day of the quarter, when you are just about to sign for the CPC upgrade. Start these discussions *now*. If the vendor refuses to be flexible, that gives you time to migrate to an alternative product/vendor before your next CPC upgrade.
- To make the administration of such agreements easier, IBM added support to SCRT in 2017 to let ISVs use SCRT, and to let you create SCRT reports showing only that vendor's products.
- If you would like this capability, **Tell Your Vendors Now**. It will take them time to develop the code and the infrastructure, so you need to give them a reasonable lead time AND you need to make it clear that you would see this as a competitive advantage for them.

Container Pricing



In July last year, IBM previewed *two* things under the Container Pricing heading.

- The FIRST is a set of infrastructure enhancements that make WLM the hub for associating workloads with pricing options, and to cascade that association through z/OS, into SCRT, and on to IBM.
 - Workloads can run in dedicated LPARs, or collocated with other workloads – *your choice*.
 - Identification of workloads is at the address space or enclave level (unlike MWP or zCAP).
 - WLM tracks CPU consumption of each 'container'.
 - Container is a logical construct that contains all work associated with that pricing option. It has **NOTHING** to do with Docker Containers, Secure Service Containers, etc.
 - The Container is used to track CPU consumption and can be used to cap the workload.
 - R4HA MSUs for each container are available using new fields in SMF Type 70 records.
 - SCRT subtracts container R4HA MSUs from 'real' R4HA MSUs.
 - Note that the MSUs reported in RMF PP Workload Activity report are 'real' MSUs, **not** R4HA MSUs. To get R4HA MSUs for the container, you must use the RMF Overview report. See *Session z101696 RMF - The Latest and Greatest* by **Karin Genther**.
- Overall effect should be roughly as if the container work was moved to another environment.

Container Pricing

- The *SECOND* part of the announcement was a set of three new *pricing options* that *exploit this new infrastructure*:
 - A DevTest option. We believe that this will be of interest to the largest set of customers of the three currently-announced offerings.
 - For more information about how this works, see our Session z104137 from Tuesday.
 - A New Application Solution option. *ONLY* for new applications. z/OS price is 50% of MzNALC price, but *all* other IBM SW used in that container must be licensed on IPLA basis.
 - Payments Solution. Based on IBM Financial Transaction Manager, the monthly charge is based on number of payments processed, not the MSUs consumed.
- In all cases, you pay for the workload as if it was in a self-contained environment - that is, regardless of whether the work runs in the peak hour or in the R4HA 'white space', you must pay for it (**think CLOUD**).
 - This means that, while attractive in some situations, it is NOT a no-brainer.
- These are just the first 3 pricing options, more are on the way.
- The latest Tuning Letter (2018 No. 1) has the first part of an article on Container Pricing.

Container Pricing

- Last week, two announcements provided:
 - DevTest enhancements ([218-324](#)):
 - Pricing Calculator (for DevTest LPARs only)
 - Better definition of DevTest
 - Another pricing metric for NewApp solution ([218-325](#)):
 - Solution Consumption License Charges (SCLC)
 - Based on total MSU consumption for the month
 - Either pay-as-you-go or baseline commitment with a discount
 - Price tiers and prices will be announced November 23rd

Cloud Pricing

- We *LOVE* the cloud...



- After all, the first cloud ran z/OS' great-great-grandfather! (If only IBM had called it 'cloud' back then....)
- *Other people LOVE* the cloud too....



- Cloud computing is projected to increase from \$67B in 2015 to \$162B in 2020 attaining a compound annual growth rate (CAGR) of 19%. [🐦](#)
- Gartner predicts the worldwide public cloud services market will grow 18% in 2017 to \$246.8B, up from \$209.2B in 2016. [🐦](#)
- 74% of Tech Chief Financial Officers (CFOs) say cloud computing will have the most measurable impact on their business in 2017. [🐦](#)

Cloud Pricing

The common perception - "The cloud is **SIMPLE** and **CHEAP!**"

A *small* sampling of the things you get to pay for when using cloud services:



Committed capacity.

Additional capacity.

Additional capacity at peak hours.

Additional capacity at peak loads.

Disk space.

Disk reads.

Disk writes.


Each application call to the database manager.

Each database manager response to the application.

Each byte that goes in and out of each server.

Placing the application and database servers in the same data center.

The *pièce de résistance* - suites of products/services to help you understand your super-simple Cloud bills!



Cloud cost containment


One popular theory is that CIOs will save money by investing in public cloud software, but that's not always the case. The fact that most CIOs leverage multiple cloud providers means enterprises are already waist-deep in complex cloud vendor management. Also, if companies leave public cloud instances running through the weekend when they don't need them, CIOs can actually spend more money than they did with on-premises solutions.

Cloud Pricing

AND, you get to select from this list of chargeable products and services for each and every one of your applications!! Is that great or what?!!

I wonder how well all of these Amazon products run when I want to move all my Cloud workload to Google because they have a Black Friday special?

Compute	Networking & Content Delivery
Amazon EC2	Amazon VPC
Amazon EC2 Auto Scaling	Amazon CloudFront
Amazon Elastic Container Registry	Elastic Load Balancing
Amazon Lightsail	
AWS Batch	Developer Tools
AWS Elastic Beanstalk	AWS CodeStar
AWS Fargate	AWS CodeCommit
AWS Lambda	AWS CodeBuild
	AWS X-Ray
Storage	AWS Tools & SDKs
Amazon Simple Storage Service (S3)	
Amazon Elastic Block Storage (EBS)	Management Tools
Amazon Elastic File System (EFS)	Amazon CloudWatch
Amazon Glacier	AWS Auto Scaling
AWS Storage Gateway	AWS CloudFormation
AWS Snowball	AWS CloudTrail
AWS Snowball Edge	AWS Config
AWS Snowmobile	AWS OpsWorks
	AWS Service Catalog
Database	AWS Systems Manager



https://aws.amazon.com/products/?nc2=h_ql_p&awsml=ql-1

Cloud Pricing

- We are not saying that enterprises should not exploit the cloud. Like everything, there are scenarios that it is perfect for. And others that it is not.
- Just don't believe all the hype. Cloud pricing is **not** as simple as you might be led to believe.
- If your company is contemplating a move to the cloud, ensure that someone does a fair, side-by-side comparison of *ALL* the costs of *ALL* your applications running in a cloud service to the known costs of your existing service.
- Don't allow yourself to be conned into comparing the cost/complexity of ONE cloud application with a z/OS environment supporting hundreds of applications.
- PS. Have you heard of 'hackers'? Are you comfortable handing over the keys to your corporate jewels to someone else?
- PPS. As far as I know, they have not released IMS for Windows (or even Linux) yet. So don't forget to budget a few \$s to **rewrite** all your mainframe applications....
- The latest Tuning Letter (2018 No. 1) has an article about this topic.

JES2/SMF

- By default, the SMF type 30 records for a job are saved in the EVENTLOG JES2 data set for each job.
- If you know how to allocate the EVENTLOG data set, you can use IEBGENER to copy the SMF records, making it SO EASY to process the SMF records for a job - no need for exits or waiting for SMF data to be offloaded...
- **Tom Wasik** kindly gave us a sample Assembler program to get us started, and our colleague **Mario Bezzi** converted that into a Rexx exec.
 - The current issue (2018 No. 1) provides a sample SDSF Rexx exec that goes into DA for a job/STC, presses enter every x seconds, and saves selected fields in a data set. If you do not subscribe to the Tuning Letter, email us for a copy.
 - The next issue of the Tuning Letter (2018 No. 2) will contain the exec and describe our experiences.

VSAM

- Prior to z/OS 2.3, CA Reclaim:
 - Defaulted to OFF at the system level (controlled via CA_RECLAIM parm in IGDSMSxx)
 - Defaulted to ON at the data class level (controlled via CA_RECLAIM attribute in data class).
- Starting with z/OS 2.3, the system-level default changes to DATACLASS. If you accept all the defaults, this has the effect of enabling CA-RECLAIM for all new data set allocations.
- There is also a new HealthCheck (VSAM_CA_RECLAIM) in z/OS 2.3 that complains if CA Reclaim is not turned on at the system level.
 - HealthCheck is rolled back to z/OS 2.1/2.2 with [OA51002](#).
- For more information, see **Glenn Wilcock's DFSMS Latest and Greatest, Session z101506**.

VSAM

z/OS 2.2 added a feature called *Dynamic Buffer Addition* for data sets open in LSR mode:

- Adds LSR buffers automatically.
- This results in IDA9990I messages and buffers being added but never removed. No way to enable/disable it at the dataset/JCL/system level.
- Controlled via BLDVRP macro.
- [OA54666](#) addresses a potential where too many buffers were being added, potentially resulting in the application running out of virtual storage. To avoid/reduce the risk of this happening, OA54666 imposes a limit of 15 additional buffers.
- Also, CICS APAR [PI92486](#) disables VSAM DBA for CICS LSRPOOLS.

The net is "DBA is not a substitute for well tuned LSRPOOLS."

zIIPs

zIIPs are CHEAP(er) than general purpose CPs. ✓

HW maintenance cost of zIIPs is lower. ✓

zIIP utilization does not count towards your R4HA. ✓

Newer versions of products often deliver increased exploitation of zIIPs. ✓

The following recent APARs add or increase zIIP exploitation:

[PI88156](#) - IMS High Performance Load - Default **OFF**

[PI88158](#) - IMS DB Reorg Expert - Default **OFF**

[PI88237](#) - Db2 Common Collector - Default **ON**

[PI89050](#) - IMS High Performance Unload - Default **OFF**

[PI91678](#) - Db2 STAP - Default **ON**

[PI92708](#) - IMS STAP - Default **OFF**

[PI92715](#) - IMS Index Builder - Default **OFF**

[PI93115](#) - Db2 High Performance Unload - Default **OFF**





COBOL Hints & Tips

- Current IBM recommendation - Migrate to COBOL V6.2 (GA 9/8/2017 which supports z14 and new Vector Packed Decimal Facility)
- Our recommendation - Migrate to COBOL V6.2 for programs run during peak R4HA (Rolling 4-Hour Average), applications that are being changed a lot, and new applications
- End of Service for COBOL V4 'might' be earlier than 2020
- 25% of customers migrating to V5/V6 have problems as a result of COBOL programs processing invalid data at run time due to compiler differences
- Cannot combine COBOL V5/V6 with OS/VS COBOL; some problems with VS COBOL II communication, so use CBT 321 program COBANAL/COBANALZ to identify load modules compiled with old versions of COBOL

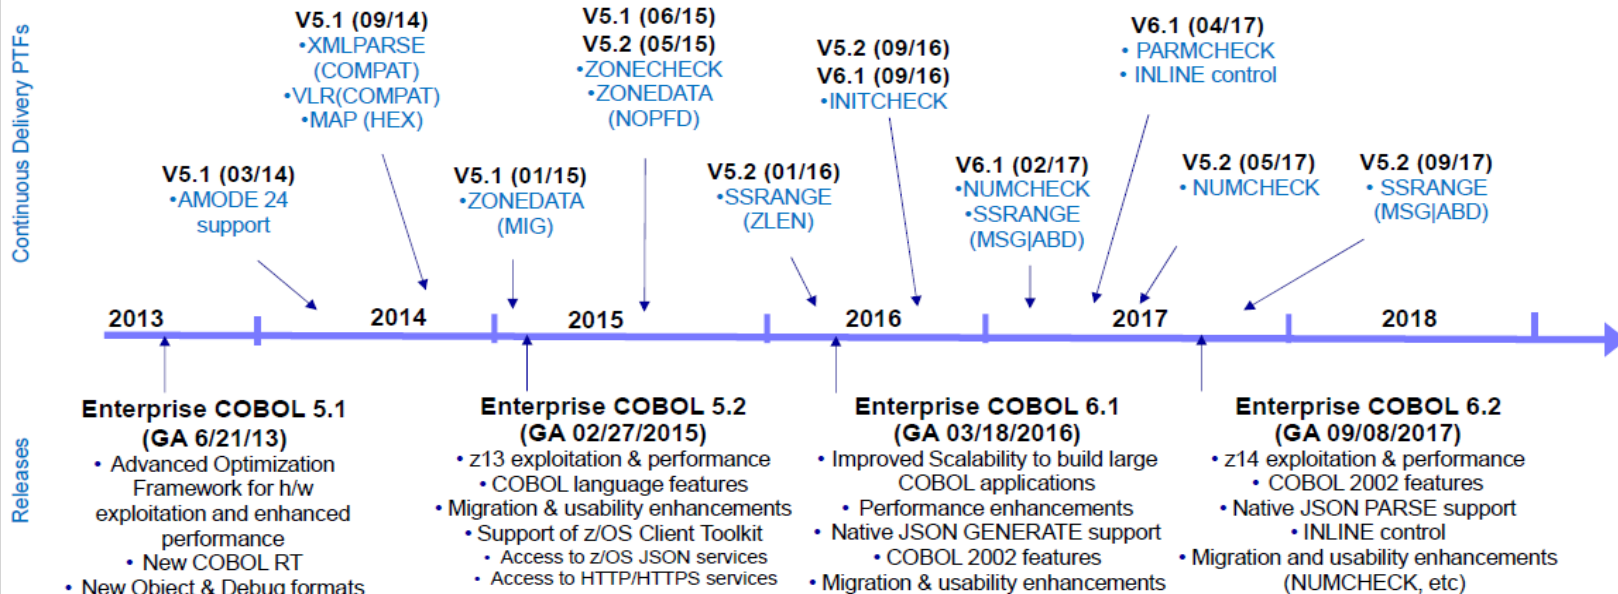
COBOL Hints & Tips

- Use new ARCH level for new hardware and recompile after moving to new machine
 - ARCH(11) for z13; ARCH(12) for z14
 - Can save 10-20% CPU time and elapsed time
 - Compile times and resources are higher
 - Don't use ARCH level higher than your D/R machine
- CICS COBOL programs may not get any benefit due to initialization overhead
 - Migrate these last because no or little benefit

COBOL Hints & Tips

- From IBM Session [22328](#):
- *To find out if users have invalid data, IBM has recommendations for migrating to COBOL V6. The first time that you compile a program:*
 1. *Compile with SSRANGE, NUMCHECK, PARMCHECK and OPT(0) for initial code changes and unit test - To find table misuse, invalid data use and invalid parameter usage - OPT(0) programs are easiest to debug, quicker compiles - Look at runtime logs for NUMCHECK, etc, error messages*
 2. *Recompile with NOSSRANGE, NONUMCHECK, NOPARMCHECK and OPT(2) plus INITCHECK for quality assurance test and production - NOSSRANGE, NONUMCHECK and NOPARMCHECK are required for good performance - OPT(2) is preferred for good performance in production - Inspect listings for INITCHECK messages*
- *Note: You may have to change to a 2-compile development process if you are not using one already*

COBOL Releases and Continuous Delivery



From SHARE Session 22341

COBOL Vector on z14

- Vector Packed Decimal Facility
 - From session SHARE 22328 and IBMTechU z100653:
 - Support for 16-byte vector registers to perform the most common packed and zoned decimal calculations is available on the z14, providing significant CPU savings
 - 4.85X faster for Unsigned Packed Decimal Add
 - 135X faster for Large Decimal Divide
 - 39X faster for Large Decimal Multiply
 - New support in CPU MF provides counts of the numbers of decimal instructions

z14 New CPU MF Counters to indicate COBOL “Modernization”

- **3 New z14 Extended Counters – See SA23-2261-04**
 - E224 - Count of floating point execution slots used for finished Binary Coded Decimal to Decimal Floating Point conversions
 - E225 – Count of floating point execution slots used for finished vector arithmetic Binary Coded Decimal instructions
 - E226 – Decimal instructions dispatched
- **Above Counters are not directly comparable to B01 (Instructions) or among each other. They could be used as an indicator of COBOL compiler “modernization”**
 - E226 – Decimal “instructions”
 - E224 – Decimal Floating Point Converted – COBOL ARCH(10 |11)
 - E225 - New z14 Vector Packed Decimal Facility and z/OS 2.3 – COBOL V6.2 ARCH(12)
- **One could identify when most Counter activity is occurring, then identify Jobs / Programs (e.g. zBNA) to investigate / re-compile for most impact**
- **See Performance examples in Back Up**

COBOL ABO


- ABO - Automatic Binary Optimizer
 - Converts load objects to more efficient code based on ARCH level (e.g. z14 is ARCH(12))
 - Can reduce run time
 - IBM says ABO takes less testing than migration to COBOL V6, but we have found that not to be the case.
 - We're looking for customers who have had good experiences; so far we've only found bad experiences (takes a lot of CPU time, costs more than it saves, output results not always equal)

COBOL

- References

- [22341](#) - *Elevating Application Performance with Latest IBM COBOL Offerings*, **Tom Ross** - great success stories
- Website - [IBM Enterprise COBOL for z/OS Migration Assistant](#), steps you through migration (requires IBM ID), has great white paper by **Tom Ross** on PDSE
- [22328](#) - *How to Take Advantage of the New COBOL v5/v6 Compilers - Migration!*, **Tom Ross** - contains COBOL Best Practices and migration steps (including using COBOL [FIXCAT feature](#))
- [22208](#) - *WSC Short Stories and Tall Tales*, by **John Burg**
- WSC Techdocs White Paper [WP102731](#) - *COBOL Applications: Techniques to Make Them Efficient*, **Kathy Walsh & Priyal H Shah**
- IBMTechU z100653 - *Increase application performance and reduce operating cost with IBM's cutting-edge COBOL offerings*, **Roland Koo**

JES2 Items of Interest



- JES2 V2R3
 - For all new items plus PTFs, see session [21805](#), *z/OS V2R3 JES2 Product Update and Latest Status* by **Tom Wasik**
 - My favorites are Email support for jobs/steps; \$D LIMITS command; dynamic scheduling group; and Init Deck Checker (will do some checking beyond syntax)
 - New JES2 SMF record 84.21 **Woohoo!!**
 - Dynamically modify SMF buffer limit with \$T SMFDEF, BUFNUM=nnn
 - Privileged space is a percent of resources held back for emergencies with normal space runs short; authorized user can logon to emergency subsystem (e.g. LOGON KYNE SUBSYS(HASP)) to clean up problems
 - Several compatability/toleration PTFs needed BEFORE starting JES2 2.3
 - **Note: z/OSMF required for JES2 Email support**

Change to IBM z/OS Support

- IBM is gradually moving support interfaces for other platforms to Salesforce-based systems.
- See https://www.ibm.com/developerworks/community/blogs/IBMElectronicSupport/entry/IBM_is_transforming_Support?lang=en and video at https://mediacenter.ibm.com/media/IBM+SupportA+Introducing+a+new+customer+portal/1_gumux5dp Not here for z yet, but coming soon to a PC near you...

Please
complete
the session
survey!

Z104139

