The Cheryl & Frank zRoadshow

Cheryl Watson (cheryl@watsonwalker.com)
Frank Kyne (frank@watsonwalker.com)
Watson & Walker, Inc.
www.watsonwalker.com
technical@watsonwalker.com
Session 16461
March 6, 2015

#SHAREorg

SHARE is an independent volunteer-run information technology association that provides **education**, professional **networking** and industry **influence**.

Welcome

- Hi, thank you for coming

- Who are we and what we do:
  - Three-person company; started in 1987
  - Frank Kyne joined us in 2014 from ITSO
  - Quarterly subscription-based newsletter –
    Cheryl Watson's Tuning Letter
    Cheryl Watson's System z CPU Chart
  - Classes on z/OS new features, WLM, performance, software pricing, chargeback, Parallel Sysplex, and high availability
  - Consulting on all the above topics
  - Software products – GoalTender and BoxScore

Agenda

- z/OSMF Workflows
- Reducing Software Costs
  - zIIPS
  - IBM Pricing Options
- Tips From Presenters
- COBOL Items of Note – Part 1
- COBOL Items of Note – Part 2
- z13 (Yes there's even more to think about!)
  - Large memory
  - SMT
- Erase on Scratch Enhancements
- Frank's baby – zPDT
- APARs
- Level 2 Problems
- Another Toy

My Favorite This Week - Workflows

- My passions over the years…
  - SMF
  - WLM
  - z/OSMF
- z/OSMF
  - Reduces sysprog time
  - Platform for all sysadmin tools in the future
  - Performance issues prior to z/OS 2.1, but now fixed with WAS Liberty Profile
  - It's critical for every z/OS 2.1 site to start exploiting z/OSMF!

z/OSMF Workflows

- z/OSMF Workflows – What Are They?
  - z/OS Project Management tool, ToDo List, Step Checklists, Migration Actions, Communication Tool
  - You can create your own step-by-step checklists, assign people to tasks, let everyone see status and dependencies
  - IBM can create workflows for migration or any other task
  - ISVs can create workflows for their own processes
  - Easy to use, easy for communication

z/OSMF Workflows

- z/OSMF Workflows – What Are They?
  - Future migrations will all use workflows
  - Most important now – z/OS Migration Workflow to perform migration actions; Marna Walle says that Migration Workflow will be kept in sync with the Migration Guide
  - Workflows have their own website: http://www-03.ibm.com/systems/z/os/zos/tools/downloads/zosmf-zos-v2r1-migration-workflow.html
  - Following slides are from Greg Daynes' session 16643 – *The New and Improved z/OSMF V2.1*

z/OSMF Workflow

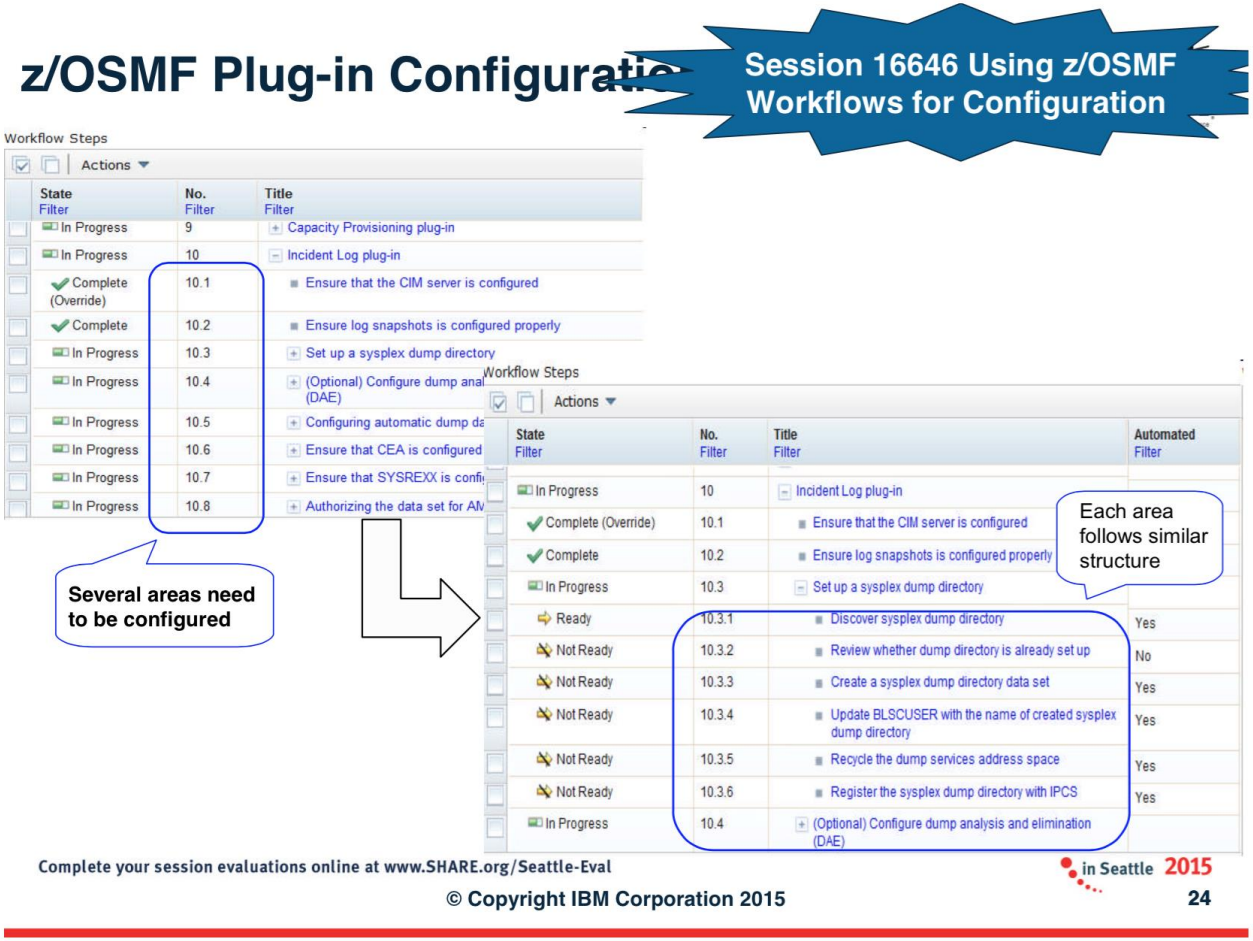

z/OSMF Workflow

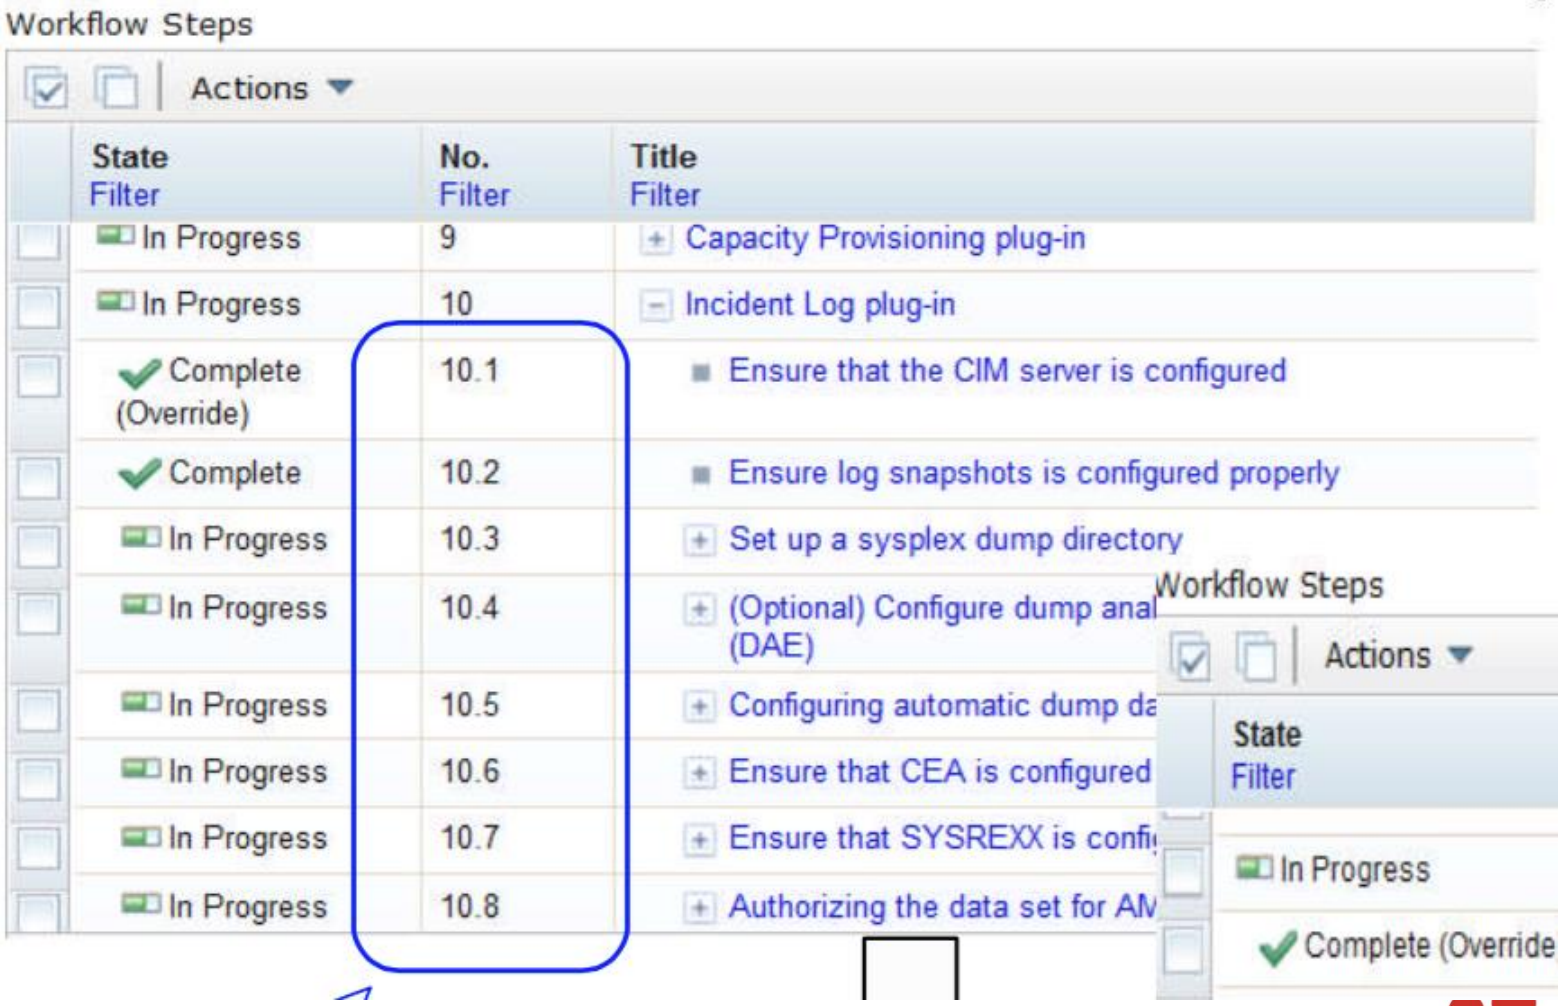

z/OSMF Workflow

In Progress	10.3	Set up a sysplex dump directory	
Ready	10.3.1	Discover sysplex dump directory	Yes
Not Ready	10.3.2	Review whether dump directory is already set up	No
Not Ready	10.3.3	Create a sysplex dump directory data set	Yes
Not Ready	10.3.4	Update BLSCUSER with the name of created sysplex dump directory	Yes
Not Ready	10.3.5	Recycle the dump services address space	Yes
Not Ready	10.3.6	Register the sysplex dump directory with IPCS	Yes
In Progress	10.4	(Optional) Configure dump analysis and elimination (DAE)	

Reducing Software Costs

From Jonathan Adams, BMC, Keynote Session 17126:

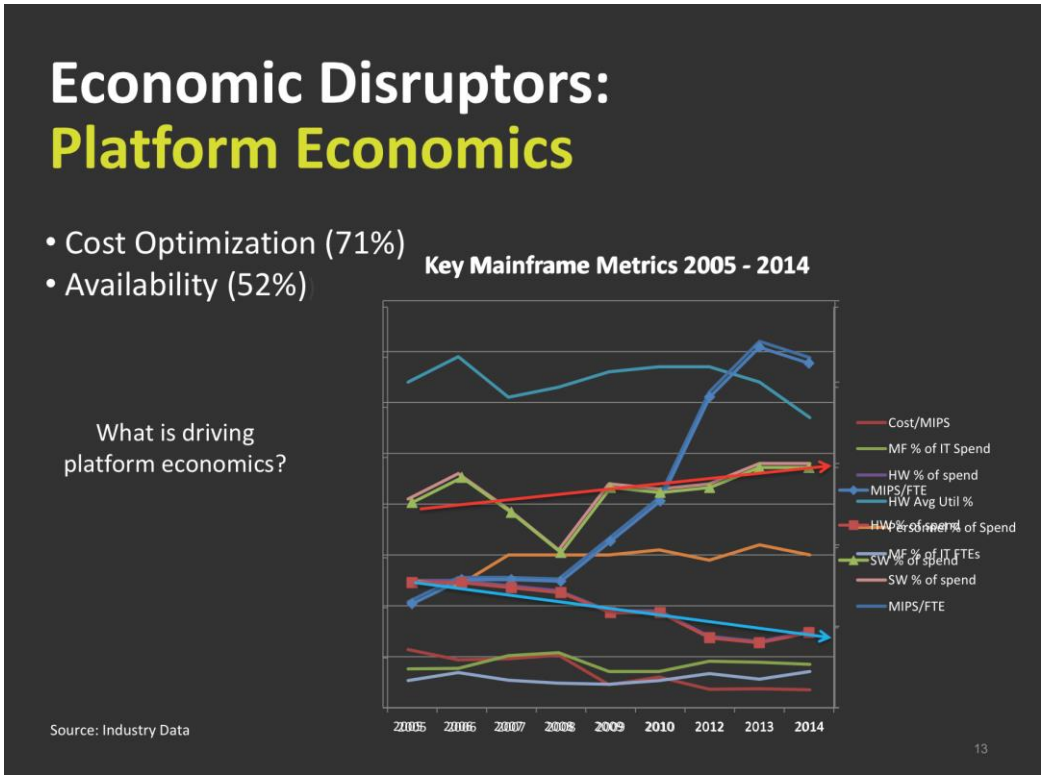

For tired (or older) eyes

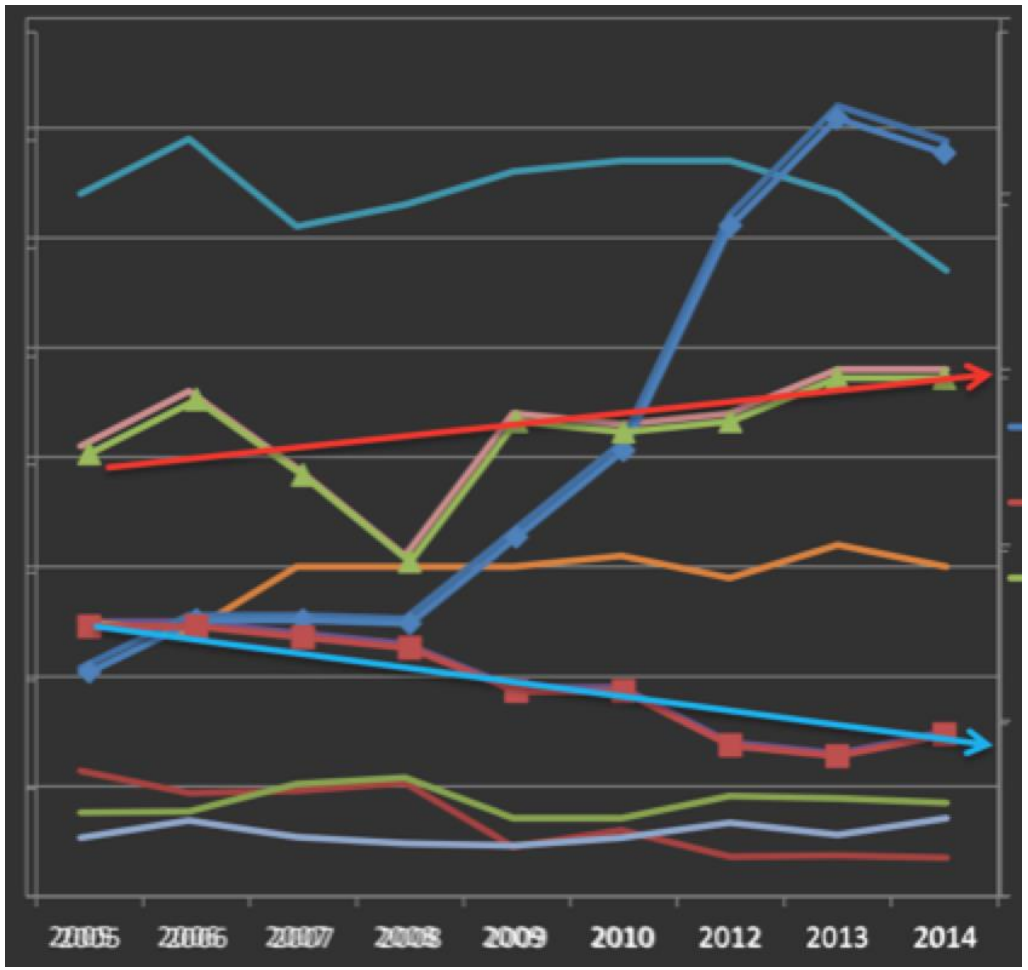

MIPS/FTE

S/W % of spend

H/W % of spend

Reducing Software Costs

- My takeaway from this graph:
  - MIPS/FTE – MIPS per Full Time Employee is going up dramatically – goodness!
    - This will continue as more automation is added
    - Those who exploit new z/OS features (e.g. z/OSMF) will achieve higher MIPS/FTE
  - 'S/W % of spend' going up and 'H/W % of spend' going down
    - How to cut S/W costs?
    - Many sessions this week on capping to reduce 4HRA
    - IBM's announcement of new pricing options
    - Exploitation of zIIPs can help
    - Use of sub-capacity CPC models can help

zIIPs – One of the biggest cost reducers

- Work run on a zIIP incurs no software costs
- z13 has no zAAPs, just zIIPs
- Greatest benefit seen in DB2, Java, WAS, MQ, z/OSMF
- But don't ignore other applications and software
- Greatest benefit – move work from CPs to zIIPs, especially during the 4HRA
- From Jeff Magdall, session 16635:
  - Uses SMT to address capacity planning around growth of zIIP-eligible work; could get up to 38% more capacity
  - ***Java SDK 8 applications1 running with z13 SMT enabled zIIP specialty engines can achieve throughput improvements of up to 50% compared to Java 7 on zEC12***

zIIPs – One of the biggest cost reducers

New Machine	U.S. Price	Avg RNI MIPS (uni)	$ per MIPS
z890*	$125,000	333	$375
z990*	$125,000	413	$303
z9-BC	$95,000	462	$206
z9-EC	$125,000	560	$223
z10-BC	$47,500	661	$72
z10-EC	$125,000	901	$139
z114	$40,000	782	$51
z196	$100,000	1202	$83
zEC12	$100,000	1514	$66
zBC12	$40,000	1064	$38
z13	$100,000	1695	$59

* - zAAPs only

Reducing Software Costs

- Recent IBM Pricing Options
  - 3/6/2014 – US Announcement 214-223 – IBM Mobile Workload Pricing (MWP)
    - Can provide 60% reduction of Mobile workload MSUs
    - Two sites at SHARE this week: one site used MWP to save $18,000/month; another site saved $38,000/month
    - Biggest problem is to determine which transactions are mobile
    - Requires new Windows-based MWRT (Mobile Workload Reporting Tool) for reporting to IBM
    - See *Cheryl Watson's Tuning Letter 2014 No. 4* for more detail

MWP

- We have all heard about Mobile Workload Pricing and most of the discussion has been about how to identify your MWP-eligible transactions.

- But there is another consideration –
  - The net effect is that some of your MSUs will have one price
  - And other MSUs (those used by MWP-eligible workloads) have a different, discounted, price

  - What does this mean if you use capping to control your software bills?

MWP

- What does your CFO care about? How many MSUs you are consuming? Or how much you pay for software each month?

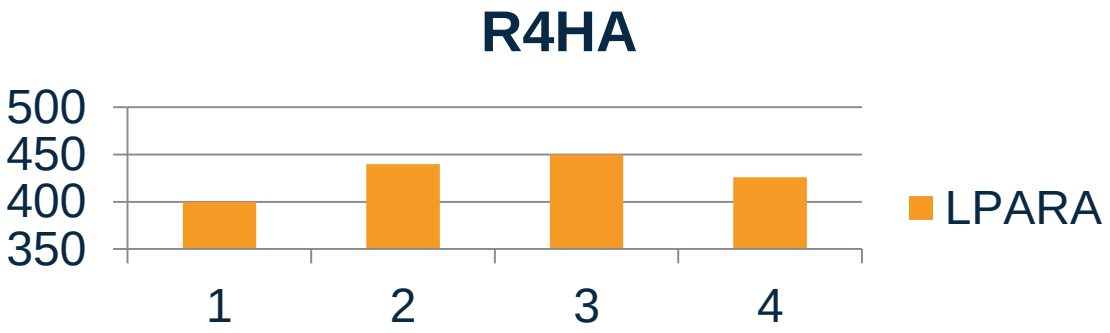

- If your MSUs cost $400/ MSU and your monthly SW budget is $180,000, you could cap the LPAR at 450 MSUs.

MWP

- What does MWP do to you?

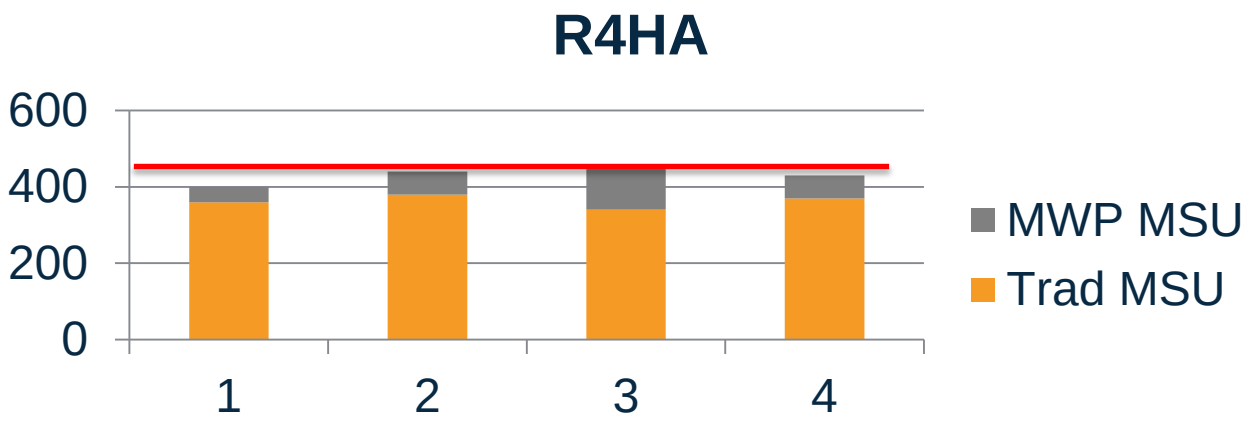

  - Peak MSUs is still 450, but 340 of those cost $400/MSU and 110 cost $160/MSU. The cost of that is 153.6K – woohoo! Let's increase the cap so we can do more work for the same price.

MWP

- Oh….

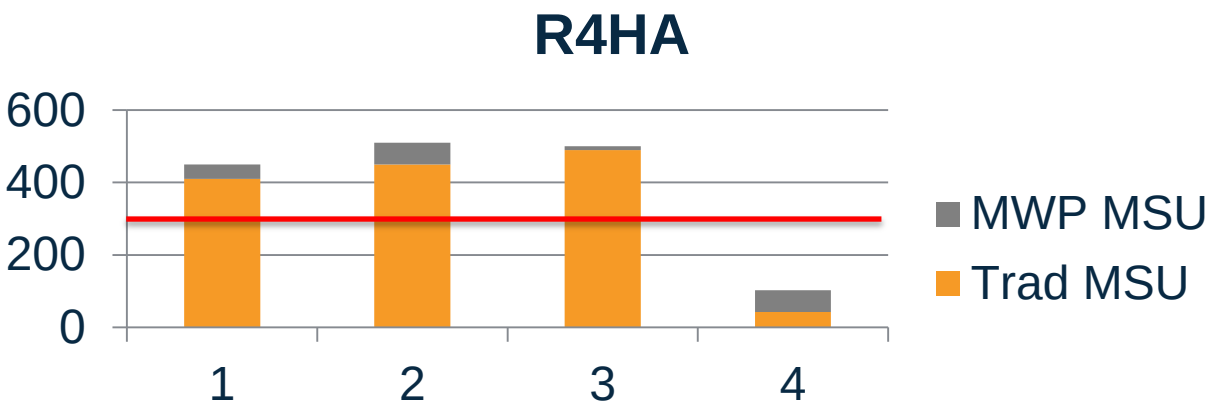

  - Peak MSUs is now 500 (because you increased the cap), but 490 of those cost $400/MSU and just 10 cost $160/MSU. The cost of that is 197.6K. Oops.
- The challenge is that your control is specified in MSUs (**NOT** traditional MSUs and MWP MSUs), but your budget is in dollars…

MWP

- The challenge is that you control your LPARs sizes based on MSUs (**NOT** separate controls for traditional MSUs and MWP MSUs), but your budget is in dollars…
- One of the attributes of mobile workloads is that their arrival rate is less predictable than traditional workloads and they tend to be more spiky.
- If you keep the old cap, you control your SW costs but you could have performance issues.
- If you increase the cap to cater for increased mobile workload there is a risk that your traditional work (batch maybe) could use all that weight.
- We still believe that MWP is a valuable option for sites with significant mobile workloads, but you need to consider how you will control it.

Reducing Software Costs

- Recent IBM Pricing Options
  - 1/14/2015 – US Announcement 215-001 – *Technology Transition Offerings for the IBM z13 offer price-performance advantages*
    - Technology Update Pricing (TUP) for z13
      - Reduces cost of certain S/W on a z13 if stand-alone or in parallel sysplex with only other z13s
      - Provides average of 5% price reduction (2.8% for 45 MSUs up to 9.8% for over 5476 MSUs)
      - Also provides reduction if z13 in parallel sysplex with zEC12/zBC12 servers (TC3); and if z13 in parallel sysplex with z196 or z114 servers (TC2)

Things that make me laugh…

- From announcement 215-001 describing TC2:

Schedule of AWLC reductions for Transition Charges for Sysplexes (TC2)

Percentage of z13, zEC12, and zBC12 Full-Capacity MSUs in an actively coupled Parallel Sysplex or Loosely Coupled Complex	Reduction in Monthly AWLC
0 - 20%	0.5%
21 - 40%	1.5%
41 - 60%	3.0%
61 - 80%	4.0%
81 - < 100%	4.5%

To find the percentage of z13, zEC12, and zBC12 full-capacity MSUs:

1. Add the total MSUs of all the servers in the sysplex or complex.
2. Divide the sum of the z13, zEC12, and zBC12 MSUs by the total MSUs in the sysplex or complex.
3. The result of the calculation must be rounded to the nearest whole percent.
   Using standard rounding rules, fractions of a percent of 0.5 or more are rounded up and fractions of a percent less than 0.5 are rounded down. Rounding examples:

   0.204999 rounds down to 20%
   0.205000 rounds up to 21%
   0.404999 rounds down to 40%
   0.405000 rounds up to 41%

4. Refer to the Schedule of AWLC reductions section to determine the applicable reduction.

Reducing Software Costs

- Recent IBM Pricing Previews (expected sometime in 2015)
  - **IBM Collocated Application Pricing (ICAP) –** Use for new application moving to z/OS; price is based on usage of new application, which is treated as if on dedicated LPAR and 'removed' from WLC of the LPAR it runs in; applicable to z114, z196, zBC12, zEC12, and z13 and certain products. ICAP reported in MWRT.
  - **Country Multiplex Pricing –** Multiplex composed of all z Systems in a country measured as if a single machine; no need to continue sham-plex to get AWLC pricing! Will have new SCRT tool. Customer needs to sign up for it and determine a baseline for future growth. This can be big! **Whoohoo!**

Tips From Presenters

- Harry Yudenfriend, Session 16896, *IBM z Systems z13 and DS8870 I/O Enhancements*
  - IBM System z I/O Exerciser
    - https://www.ibm.com/services/forms/preLogin.do?source=swg-beta- ibmioexzos
      New tool made available March 4, 2014 to verify quality of the cable connections before running z/OS production work
    - Runs in a stand-alone LPAR or z/VM Guest Machine
    - Tests all the FICON devices available to that partition via the IOCDS
  - Rest of handout has some impressive performance I/O improvements obtained by exploiting new functions

Tips From Presenters

- John Burg, Session 16803, *2015 CPU MF Update*
  - Reiterated importance of turning on HIS and collecting SMF 113 records; customer data needed to fine tune capacity estimates for zPCR:

Looking for “Volunteers”

(3 days, 24 hours/day, SMF 30s, 70s, 72s, 99 subtype 14s,113s per LPAR)
“Before z196 / zEC12” and “After z13”
Production partitions preferred

If interested send note to jpburg@us.ibm.com,
No deliverable will be returned
Benefit: Opportunity to ensure your data is used to influence analysis

Tips From Presenters

- Us too!
- Watson & Walker is looking for same type of data, and will provide reports from our various tools

Looking for “Volunteers”

(3 days, 24 hours/day,

SMF 30s, 70s, 72s, 89, 90, 99 subtype 14s,113s per LPAR)

“Before z196 / zEC12” and “After z13”

“Before z196” and “After zEC12”

Production partitions preferred

If interested send note to technical@watsonwalker.com.

We’ll be analyzing a variety of SMF data fields

and will provide you with any reports we create.

Tips From Presenters

- Before z/OS 2.2, you MUST migrate from Domino HTTP Server to Apache HTTP Server
- Rumor is that this is NOT an easy task
- But there is help:
  - IBM Redpaper – REDP-4987-01 – *IBM HTTP Server on z/OS: Migrating from Domino-powered to Apache-powered* (Newly published 01 February 2015); authored by Edward McCarthy
  - SHARE Session 16928 – *HIS Apache for z/OS – How to Implement and Exploit* by Edward McCarthy
  - Thanks Edward!
-

COBOL

- What if you no longer have the source code and can't migrate to COBOL V5 with its better performance?
  - Two options: Source Recovery Company & IBM COBOL Binary Optimizer
- Source Recovery Company
  - Recommended by Tom Ross, "Captain COBOL"
  - www.source-recovery.com
  - Can produce COBOL source from COBOL executable

COBOL

- IBM COBOL Binary Optimizer
  - Currently a research prototype - https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W70542409ccc8_47fa_bcaa_4c2dcee9c78a
  - Download tool, pick a CPU-intensive COBOL V3 or V4 module, run it, optimize it, run it again – send IBM results
  - Find CPU-Intensive programs with IBM CPU Measurement Facility - http://www-03.ibm.com/systems/z/os/zos/features/unix/tools/hisreport.html
  - Reports of 20% to 40% CPU time savings – Yowza!

COBOL

- COBOL V5 Pricing
  - Use 90-day free trial (can use it for both 5.1 and again for 5.2)
  - Can use Single Version Charging (SVC) for 12 months
  - This isn't long enough for most sites
  - If you need longer, press your IBM team
  - There are customers who have gotten extensions by showing cause
  - Longer SVC term is under consideration by IBM – make your voice heard!

COBOL

- COBOL V5 User Experience, Brian Peterson, Session 16710
  - Biggest COBOL conversion in 30 years
    - He considered OS/VS COBOL to VS COBOL II in mid-1980s to be large conversion (100 on “1-100 point” scale); Enterprise COBOL V3 was medium (10); Enterprise COBOL V4 was small (2)
    - COBOL V5 is significant – 20 on “1-100 point” scale
  - Results
    - Several thousand COBOL V5 in production
    - Removing COBOL V4 will end up taking over a year
    - Looking forward to COBOL 5.2 with new RULES feature and no SIZE parm

COBOL

- COBOL V5 User Experience, Brian Peterson, Session 16710 (cont.)
  - Winning techniques
    - Recommended a “Convert at Change” strategy
    - Created team of sysprogs (compiler, zOS, developer tools), performance analysts, source code lifecycle management, and application reps) meeting weekly with group mailbox
    - Work closely with vendors
    - Keep current with maintenance (not RSU-only) – http://www-01.ibm.com/support/docview.wss?uid=swg27041164

COBOL V5 Performance

- One of our customers noticed that the elapsed time of one of their monthly batch jobs increased by nearly 4 hours when they moved to COBOL V5.
- Further investigation showed that the number of I/Os issued by the job had increased from about 30,000 to 38,000,000.
- What's going on?

COBOL V5 Performance

- COBOL V5 *requires* that its application load libraries are PDSEs.
  - However they had already moved to PDSEs with COBOL V4 and not had any performance issues.
- To avoid the need to pre-link your modules, ALL COBOL V5 program objects contain ‘deferred segments’.
  - So do most C/C++ programs, and PL/I programs compiled with the RENT compiler option.
- Today, LLA will not place programs with deferred segments in VLF.

COBOL V5 Performance

- This customer uses LLA to manage their application load libraries, so when they moved to COBOL V5 (from V4), they lost the benefit of caching in VLF (because of COBOL V5's use of deferred segments).
- *They also didn't have the PDSE1 Hiperspace enabled.* So their programs went from being cached in VLF to not being cached at all.
  - Default is that PDSE Hiperspaces are NOT enabled.
- As a result, each of the 38M I/Os to load the program took .4ms, compared to .04 ms to load them from VLF.

- So they contacted us and we spoke to **Tom Reed**, **Peter Relson**, and **Tom Ross** in IBM….

COBOL V5 Performance

- Lessons:
  - Make sure that you enable the PDSE and PDSE1 Hiperspaces (in IGDSMSxx) BEFORE you start your migration to COBOL V5 **AND** that PDSE HIPER APAR OA46328 is installed.
    - Customer subsequently did this and said that performance recovered to near-pre-COBOL V5 levels.
  - We also recommend that you ensure that you have SMSPDSE1 set up. This, and enabling PDSE Hiperspace, requires an IPL.
    - PDSE1 Hiperspace can be enabled without an IPL if SMSPDSE1 is already set up.

COBOL V5 Performance

- Lessons:
  - Apply the PTF for APAR OA45127 when it is available – this z/OS 2.2 enhancement is rolled back to 1.13 and 2.1 and changes LLA so that program objects with deferred segments ARE eligible for VLF.
  - If you do not use LLA for your application load libraries, you should consider doing so.
  - For more information, see:
    - Cheryl's List 181 (http://www.watsonwalker.com/clist181.html)
    - Cheryl's List 182 (http://www.watsonwalker.com/Clist182.pdf)
    - IBM Technote *An Overview of Hiperspace Caching for PDSE*

z13 - It's arrived!

- The latest addition to the IBM z Systems family was announced in January 2015

IBM z13

- IBM z13 has all the improved bells and whistles we have come to know and love:
  - Faster channels (16Gb FICON)
  - More engines (up from 101 to 141)
  - More MIPS
  - More MIPS per CP
  - More LPARs (up to 85)
  - More channel subsystems (up to 6)
  - More subchannel sets (up to 4)
  - New Coupling Links (Integrated Coupling Adapter)
  - IBM zAware extended to support analysis for Linux as well as z/OS
  - New Crypto 5S

IBM z13

- Plus, some real game changers:
  - *Significantly* more memory – max up from 3TB to 10TB AND significantly lower memory pricing.

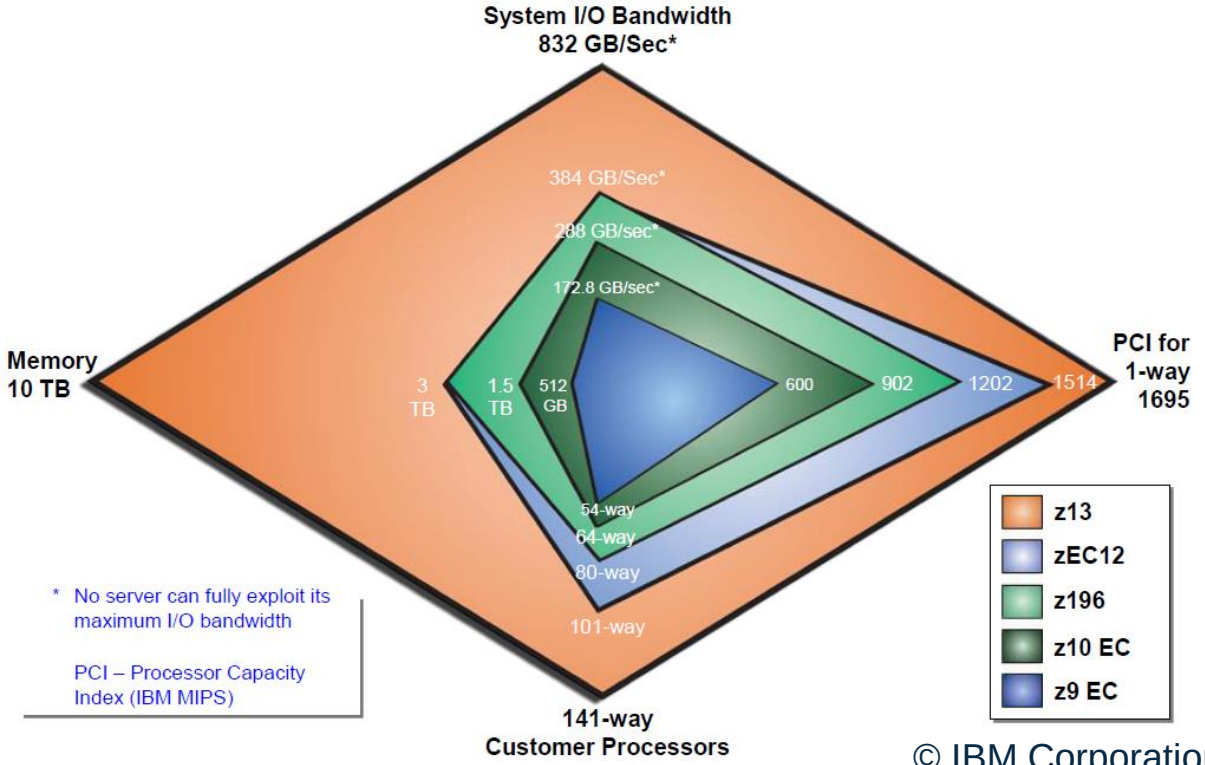

© IBM Corporation

What is Driving z13?

New marketplace dynamics will drive hyper growth opportunity for the IBM Mainframe

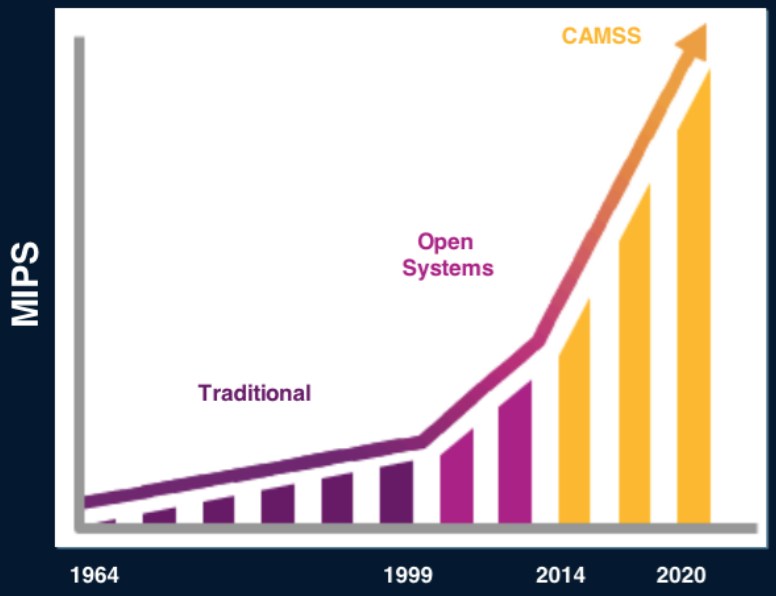

Traditional 16M MIPS, 10% CAGR
1964–2014
- Batch
- General Ledger
- Transaction Systems
- Client Databases
- Accounts payable / receivable
- Inventory, CRM, ERP

Linux & Java +14M MIPS, 28% CAGR
1999–2014
- Server Consolidation
- Oracle Consolidation
- Early Private Clouds
- Email
- Java, Web & eCommerce

CAMSS
2015–2020
- On/Off Premise, Hybrid Cloud
- Big Data & Analytics
- Enterprise Mobile Apps
- Security solutions

IBM z13

- Large memory
  - There is nothing really new in the world…. Does anyone remember Data in Memory (in vogue in the 1980s)?
    - Data In Memory was an attempt to reduce the impact of diverging CPU and I/O speeds by keeping more data in memory….
  - Large memory is a way to:
    - Allow customers to continue increasing workload volumes and complexity even though increases in CPU engine speed are decreasing.
    - Enable applications that would not otherwise be possible – in-transaction fraud detection, for example.
    - Large memory enables these by letting you put more data closer to the cores.

IBM z13

- Large memory
  - IBM is really trying to encourage z customers to exploit these capabilities, so the next time you are changing technology, ask for prices for various memory sizes, including a lot more than you might think you need…
  - IBM doesn't publish H/W prices, but we hear that you can get **three times the amount of memory for the same price** (and there are special deals!)

IBM z13

- Simultaneous MultiThreading
  - This IS completely new, to the mainframe world at least…
  - We have always had one stream of instructions on a core (CPU) at one time. So the speed at which work would get processed was pretty consistent, with the effectiveness of processor cache being the main differentiator between the performance seen by one workload type versus another.
  - You will soon look back fondly to those simpler days….

IBM z13

- Simultaneous MultiThreading
  - Hello more capacity, goodbye repeatability…
- Harv Emery had an excellent analogy of the impact of SMT:
  - Each core can now be set up as one ‘lane’ or two ‘lanes’.
  - If one, the speed will be higher, but with two lanes, even though the speed is lower, you can *potentially* get more cars down the road (capacity).

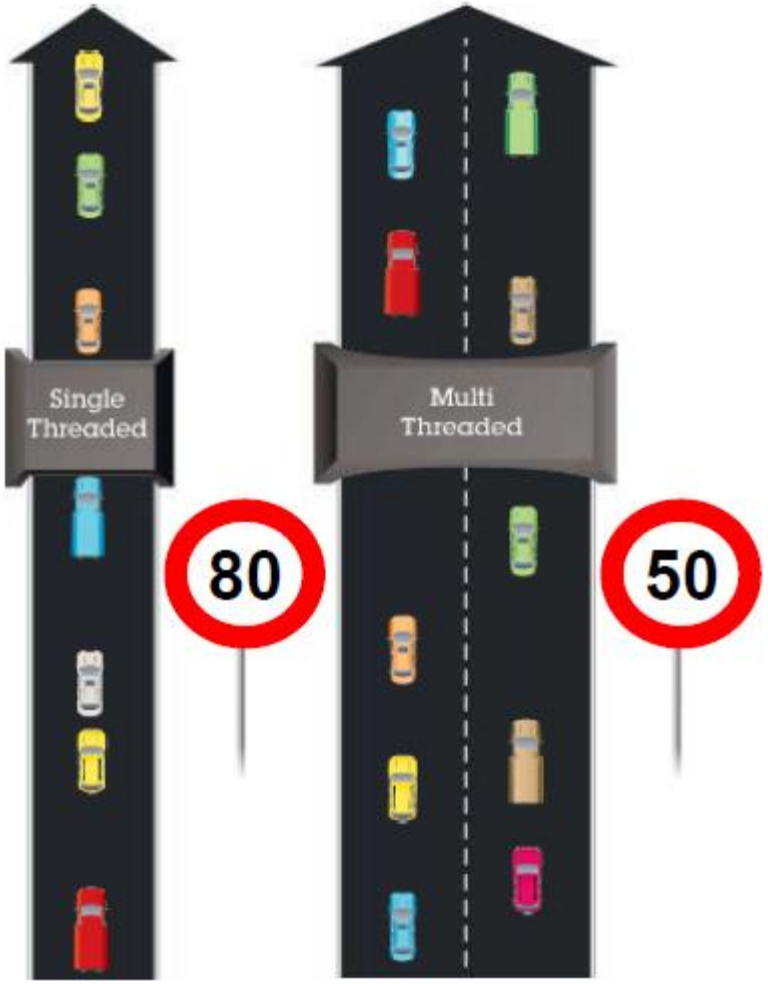

IBM z13

- What Harv's picture is missing is intermittent roadworks, where the two lanes narrow to one lane.
  - Just like a real highway, those roadworks only impact you if two cars want to use that lane at the same time.
- Based on IBM's experience so far, with artificial workloads, SMT2 (2 lanes) delivers between 10% and 40% more capacity than SMT1 (1 lane).
  - However, this does not mean that one workload type consistently gets 10% more throughput, and a different type gets 40% more throughput.
  - Just like a real highway, the capacity depends on how often 2 cars reach the roadworks at the same time. So maybe a given CICS workload will get 10% more capacity now, 40% more 2 minutes from now, and 23% more 5 minutes after that.
  - WE DON'T KNOW YET.

IBM z13

- What we DO know:
  - In an environment using SMT2, performance and capacity are no longer the same things.
  - IBM has wisely only delivered SMT2 on zIIP (and IFL) PUs at the moment. This allows us all to get more experience in a gradual manner.
    - AND you can turn SMT2 on and off dynamically using a SET command.
  - This is the way the industry is going – it was just a matter of time before IBM was forced to do this for z (they did it for p Series years ago).
  - Most SMT implementations are semi-transparent to the operating system. For SMT on z, the IBM hardware, hypervisor, and operating system teams have been working together on this from the start. So IBM is hoping that the implementation of SMT on z will be more effective than SMT on other platforms.

IBM z13

- What else do we know:
  - The cache effectiveness of a workload will still be a major factor in the capacity you can drive from a given configuration – maybe even more so than in the past because of the impact of a cache miss on co-resident threads AND the new memory structure in z13. So make sure you are collecting those Type 113 records.
  - If you have been complaining (or receiving complaints) about CPU variability before z13, you ain't seen nothing yet!
  - There are a LOAD of new terms for us to get used to – CORE, MT1 Equivalent Time, Capacity Factor (CF), Maximum Capacity Factor (mCF), Average Thread Density, Core busy time, Productivity…. AND the meaning of some terms we are used to might change. My brain hurts already….
    - Plus new SMF fields, new RMF report layouts

RMF CPU report in SMT2 mode

---CPU---		---------- TIME % ----------				--- MT % ----		LOG PROC		--I/O INTERRUPTS--	
NUM	TYPE	ONLINE	LPAR BUSY	MVS BUSY	PARKED	PROD	UTIL	SHARE %		RATE	% VIA TPI
0	CP	100.00	60.59	60.46	0.00	100.00	60.59	100.0	HIGH	1888	23.72
1	CP	100.00	70.30	70.30	0.00	100.00	70.30	52.9	MED	974.7	9.62
2	CP	100.00	26.12	38.46	32.14	100.00	26.12	0.0	LOW	0.00	0.00
3	CP	100.00	13.14	33.27	60.52	100.00	13.14	0.0	LOW	0.00	0.00
4	CP	100.00	0.00	-----	100.00	100.00	0.00	0.0	LOW	0.00	0.00
5	CP	100.00	22.56	22.54	0.00	100.00	22.56	0.0	LOW	0.00	0.00
6	CP	100.00	0.00	-----	100.00	100.00	0.00	0.0	LOW	0.00	0.00
….											
TOTAL/AVERAGE			13.76	47.27		100.00	13.76	152.9		2863	18.92
E	IIP	100.00	0.23	0.19	0.00	100.00	0.23	100.0	HIGH		
				0.10	0.00						
F	IIP	100.00	0.07	0.06	0.00	100.00	0.07	100.0	HIGH		
				0.03	0.00						
TOTAL/AVERAGE			0.15	0.09		100.00	0.15	200.0			

------------ MULTI-THREADING ANALYSIS ----------				
CPU TYPE	MODE	MAX CF	CF	AVG TD
CP	1	1.000	1.000	1.000
IIP	2	1.000	1.000	1.240

SMT Summary

- This is not going away, and it is going to be quite different to everything we have been used to:
  - How is IBM going to size new CPCs?
  - **How are you going to handle this in your chargeback systems?**
  - How will this impact your software bill when/if IBM rolls SMT2 out for general purpose CPs?
  - How will you do capacity planning if the capacity you are consuming can change dramatically from minute to minute based on how well the transactions are interacting at the thread level at that instant? Maybe, with the volume of work on z, the peaks and valleys will get smoothed out and this won't be an issue?

  - All good things to meditate about on your flight home.. We will expect you to have come up with answers by Orlando.

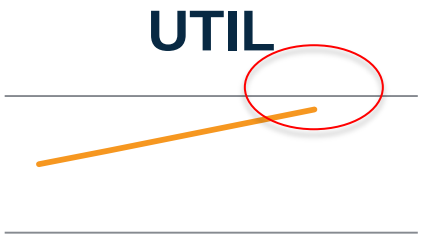

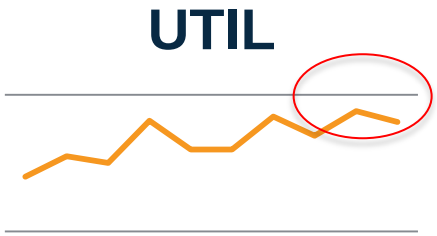

IBM z13 key dates

- January 14 – Announce
- Feb 27 – CF Sizer support for CF Level 20
- March 9 – General availability
- March 14 – zVM 6.3 support for SMT
- April 14 – TKE 8.0 LIC & related functions, AND the Universal Lift Tool Upgrade Kit for zEC12/zBC12
- June 26 – 256 Coupling CHPIDs, STP enhancements, FCP enhancements, improvements to zHPF long distance performance, zAware Linux support
- Sept 25 – FICON Dynamic Routing, Forward Error Correction for FICON 16Gb, SAN Fabric I/O Priority
- For LOADS more info, see Harv Emery's two excellent sessions on z13 (**also available on SHARE Live**):
  - 16704: The New IBM z13 Part 1
  - 16459: The New IBM z13 Part 2

Erase on Scratch Enhancements

- When a data set is allocated, data within that data set is protected by your security software.
- But what happens when that data set is deleted or part of its space is released?
  - A new data set allocation will overlay that space. Access to the tracks that belong to the new data is controlled by the security profile of the *new* data set. However, those tracks could still contain residual data from the *first* data set.
  - For a long time, this was an exposure that you had no way to protect yourself from.
- To address this issue, IBM introduced the Erase on Scratch function way back in RACF 1.7. It gives you the *option* to say that residual data should be overwritten when a data set is deleted or space is released.

Erase on Scratch enhancements

- Unfortunately, the performance of Erase on Scratch was so dreadful that I don't know anyone that used it at that time (this was before hackers were invented!).
- Things improved with IBM RVA, because it provided a function (DDSR) where the scratch processing could be driven by the control unit (upon request from DFP).
  - But it still only scratched one track per request from DFP.

Erase on Scratch enhancements

- z/OS 2.1 enhanced this processing to delete up to 255 tracks per request.
  - This happens *automatically* if the storage subsystem supports this capability. EXCEPT for PPRC primary devices – they are still scratched one track at a time.
- APARs OA43693 and OA46511 provide a new DEVSUPxx keyword (EOSV2) that let's you turn this enhancement on for PPRC primaries as well (*and the APAR text provides info about the DASD microcode levels you need to exploit this capability*).

zPDT

- What is zPDT?
- Emulator to let you run z Systems operating systems under Linux on a PC.
  - Supports most, but not all, mainframe functions.
- Costs somewhat less than a z13
  - AND it fits more easily in your backpack.
- Supports from 1 up to many users.
- Potential uses:
  - ISV development
  - Application development
  - In-house operator and sysprog training

zPDT

- zPDT forms the basis of the RD&T offering from IBM.
  - zPDT is just the emulator, it doesn't include the operating system.
  - The operating system(s) and all other z/OS software is delivered in the RD&T offering.
  - ISV vendors can get zPDT and the ADCD from IBM or ITC
- For more info, see Tom Conley's and Stephen Norris' excellent user experience presentation.
- Combine zPDT and the Mocha TN3270 emulator for Android, and you will be the envy of all your friends and neighbors

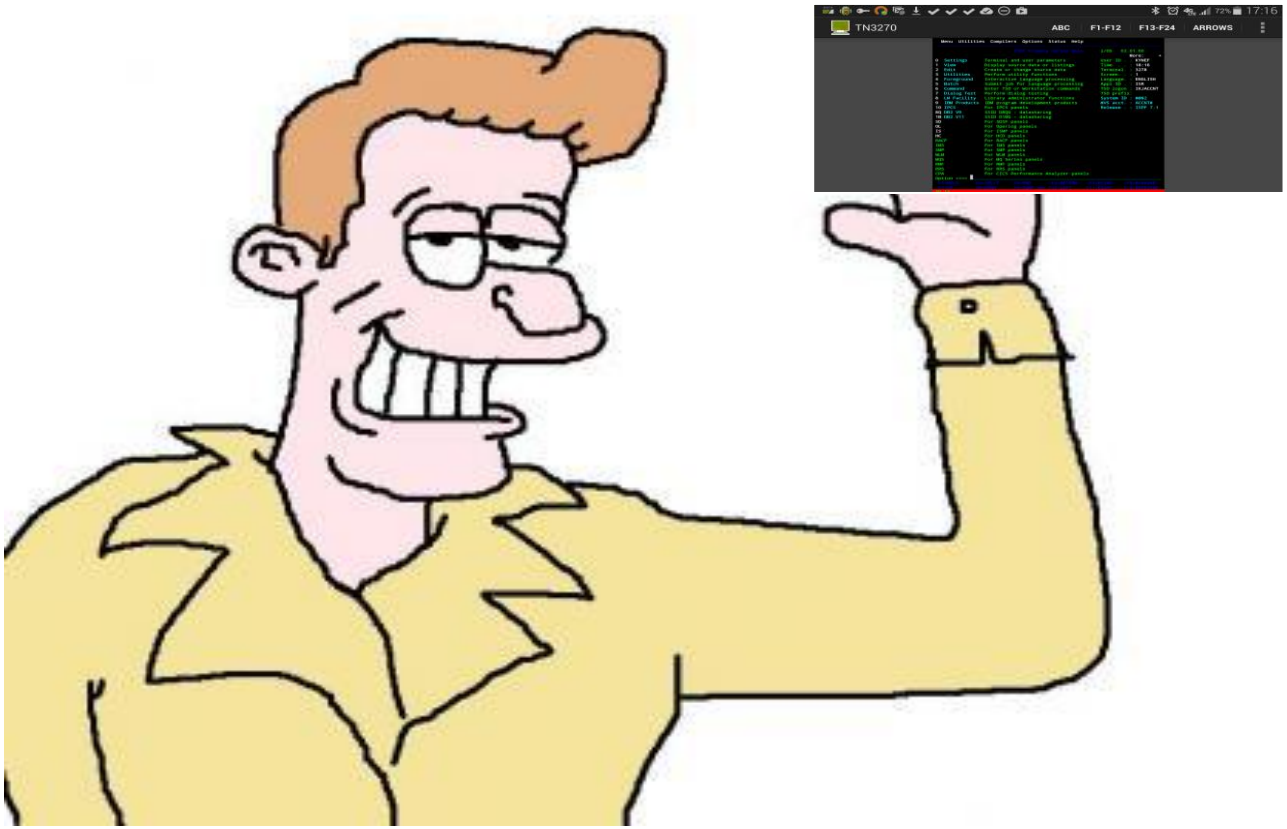

APARs I have known and loved

- This is a list of recent APARs that you HAVE to install on your systems
  - WLM: OA43622 (SMT support), OA47021 (z13 support), OA44504 and OA46396 (multiple RSM-related changes for large memory).
    - See Horst's presentation for more details.
  - RMF: OA44101 (SMT support), OA44524 (PCIE enhancements on z13), OA44502 (Support for ICA links)
    - For more information on all these enhancements, see Peter Muench's session 16816 RMF Latest & Greatest presentation.
  - zEDC: OA46498 Data Loss HIPER APAR if you are using zEDC for SMF logstreams.
  - XRC: OA43453 HyperPav XRC Workload-Based Write Pacing.
  - DB2/HyperSwap: PI25747 and OA45125 Enable zHyperWrite for DB2 (V10 & later)
  - MWP and IBM Transaction Analysis Workbench: PI29291 New Function To Support Mobile Workload Pricing.

Problems I have known and avoided

- Level 2 lessons
  - Check out the latest update to the SAD Best Practices white paper. And perform some practice stand alone dumps, including sending the dump to IBM, to ensure that you don't hit unexpected delays in case you need to send a real dump to IBM.
  - See handouts from Patty Little and John Shebey's *Level 2 Insight: Common z/OS Problems You Can Avoid* session.

Problems I have known and avoided

- Level 2 lessons
  - Poor performance with zFS 1.4 large (>10K objects) directories.
    - Implement Fast Lookup Cache
    - Download largedir.pl Perl script from Tools & Toys web site to identify large zFSs.
    - Implement zFS 1.5 if all systems in plex are running z/OS 2.1 or later.
  - zFS 1.5 flies in 2.1! If you haven't converted off HFS, NOW is the time!

One can never have too many tools

- If you think it would be helpful to understand the relationship between logical and physical CPs today, just wait till you get your z13!
- Luckily for you, IBM has a cool new tool to do this for you, called the WLM Topology Report. It processes SMF Type 99.14 records and creates a CSV file that can then be viewed on your PC. Works on CPCs back as far as z10 and as recent as z13.
  - See Horst Sinram's WLM Update session and John Burg's session 16803 for more info.
  - SMF 99.14
    - Gives physical topology of zEC12 (Book/Chip) or z13 (Drawer/Node/Chip)
    - Written every 5 minutes or when performance changed

Thank you, Thank You, THANK YOU!

Thank you for coming, have a safe trip home, and we'll see you in Orlando!

Complete your session evaluations online at www.SHARE.org/Seattle-Eval