

Cheryl's Hot Flashes #26

Cheryl Watsons
Watson & Walker, Inc
www.watsonwalker.com - home of Cheryl Watson's
Tuning Letter, CPU Charts, BoxScore and GoalTender

August 12, 2011
Session 09588

Agenda

- Introduction
- From our Tuning Letter
 - SMF Enhancements
 - IBM Resiliency
 - New Function APARs
 - Gotcha APARs!
 - Great Docs
- z/OSMF
- zEnterprise Issues
 - z114s
 - Specialty Engine Costs
 - zEnterprise Positioning
 - Power Saving Mode
- At This SHARE
- From Last SHARE



Introduction

- Currently
 - Author of *Cheryl Watson's Tuning Letter* (40-60 pages six times a year) and *Cheryl Watson's CPU Charts*
 - Author of free email – *Cheryl's List* (sign up on website)
 - Developer of two Watson & Walker's software products, *BoxScore* and *GoalTender*
 - Long-time SHARE contributor (ribbon wearer since 1978)
 - CMG past director/contributor
 - CMG A. A. Michelson Award winner
 - zJournal Mainframe Hall of Fame
 - Presents “Hot Flashes” at every SHARE to talk about the things I'm passionate about

Introduction

- History
 - 1965 – Math & physics major at Portland State; worked at Consolidated Freightways, wrote Autocoder on 1401 and 7010; installed MFT in 1966
 - 1967-1982 – Several software companies, Amdahl, and EDS (training, CICS admin, performance and capacity)
 - 1982-1986 – Morino Associates (England, Germany, Virginia)
 - 1986 – Met Tom Walker, partner and future husband, started Watson & Walker, Inc. as a training company (taught SMF, RMF, performance, WLM, capacity planning until 1999)
 - 1991 – Started *Cheryl Watson's Tuning Letter*
 - 1995 – Created WLM QuickStart Policy
 - **Never met an SMF record I didn't like!**

From our Tuning Letter – SMF Enhancements

- SMF Enhancements
 - IEFU29L exit for switching SMF when using logger is not available in SYS1.SAMPLIB. There is an MVSE requirement that it be made available (SSMVSE10010).
 - Until then, look at Redbook [SG24-7919-00](#) (*SMF Logstream Mode: Optimizing the New Paradigm*), in the Additional Materials section.

From our Tuning Letter – SMF Enhancements

- SMF Enhancements
 - The IFASMFDL and IFASMFDL SMF dump programs do not show the SMF record subtypes, so you really can't tell which subtypes are responsible for the largest amount of data. There is an EWCP requirement that the dump programs provide this information (SSEWCP08001).
 - Until then, look at Redbook [SG24-7919-00](#) (*SMF Logstream Mode: Optimizing the New Paradigm*), in the Additional Materials section. See Appendix E.

From our Tuning Letter – SMF Enhancements

- Figure 2 – Sample IFASMFDP report from Redbook SC24-7919-00

DETAILED ACTIVITY REPORT FOR OUTPUT DDNAME: SMFOUT

RECORD TYPE	RECORD SUBTYP	SYSTEM ID	RECORDS WRITTEN	BYTES WRITTEN	RECORD MINLEN	RECORD MAXLEN	START DATE	START TIME	END DATE	END TIME
2	N/A	#@\$A	1	18	18	18	10.210	17:23:43:50	10.210	17:23:43.50
3	N/A	#@\$A	1	18	18	18	10.210	17:23:51:17	10.210	17:23:51.17
14	N/A	#@\$A	35	12404	344	608	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
15	N/A	#@\$A	37	12756	344	372	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
17	N/A	#@\$A	1	100	100	100	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
23	N/A	#@\$A	13	3354	258	258	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
30	1	#@\$A	3	1221	400	413	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
30	2	#@\$A	331	450861	1119	3925	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
30	3	#@\$A	5	6361	1209	1435	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
30	4	#@\$A	5	8341	1119	2109	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
30	5	#@\$A	5	8393	1119	2122	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
30	6	#@\$A	99	107613	1087	1087	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX
60	N/A	#@\$A	1724	942538	338	547	10.210	15:XX:XX:XX	10.210	15:XX:XX.XX

From our Tuning Letter – IBM & Japan Earthquakes

- [Timothy Sipples](#) from IBM Singapore posted an item on *The Mainframe Blog* – mainframe.typepad.com/blog/2011/03/earthquake-in-japan-ibm-machines-kept-running.html.
- All of these IBM servers kept running, and there were no service interruptions.
- Now that's impressive!
- The photos on the next two slides are from the IBM data center in Tokyo and are Copyright 2011 The Mainframe Blog (mainframe.typepad.com).

From our Tuning Letter – IBM & Japan Earthquakes



From our Tuning Letter – IBM & Japan Earthquakes



From our Tuning Letter – New Function APARs

- [OA34579](#) (z/OS 1.10-1.13, 22Jun2011) – *New Function – SETXCF Modify Command Option to Disable Coupling Facility Structure.* Also provides performance tips when altering CF structures.
- [OA35146](#) (z/OS 1.10-1.12, 18May2011) – *New Function – Allow Non-client Preemptable SRBs to Join/Leave an Enclave After it has Begun.* This is a z/OS 1.13 enhancement rolled back that is used by DB2 V8-V10.
- [OA35675](#) (RMF 1.12, 27May2011) – *New Function – RMF Enhances its CPU Reporting Capabilities for Systems Running as z/VM Guests.* This lets your z/OS Partition Data Report show CPU time used by z/VM.

From our Tuning Letter – Caution APARs

- [OA36354](#) (z/OS 1.12, OPENed 29Apr2011) – *Increased CPU Activity After Migrating to z/OS 1.12 DB2 Create Index IGG029CM Calling IFFDDSRV in SYSTRACE*. Although the APAR indicates that the CREATE INDEX utility causes the CPU increase, it's actually the dropping of a DB2 object that causes it. Thanks to [Brian Scott](#), [Mike Draggs](#) & [Ron Johnson](#) from Data-Tronics for the info. After installing the ++FIX, the CPU time for DB2PDBM1 went from 19988 seconds per day to 10074 seconds per day (back to z/OS 1.10 levels). **WOW!!!**
- DB2 Level 2 also recommended these APARs for any DB2 shop going to z/OS 1.12 (or higher):
 - [PM17542](#) (DB2 V8-V9, 14Oct2010) - *DB2 Performance Improvements*. Please note that the APAR describes an additional step that is needed after applying the APAR. You need to add to the ALLOCxx parmlib member a parameter of SYSTEM MEMDSENQMGMT(ENABLE).
 - [OA36101](#) (z/OS 1.12-1.13, 18May2011) – *Overlay of High Private with a USB Address*. This corrects a problem with the z/OS 1.12 enhancement used by PM17542.

From our Tuning Letter – Caution APARs

- [PM28626](#) (DB2 V8-V10, 8Apr2011) - *Correction of DRDA Using TCP/IP Execution Variation and Handling of Enclave Control Structure Anomalies.* After the application of PTFs for [PM12256](#) (DB2 V8-V9, 8May2010 - *DRDA Performance Improvement Using TCP/IP*), users found that
 1. zIIP utilization levels became more variable,
 2. execution time of DRDA threads with multiple SQL statements became more variable,
 3. performance of a single DRDA long running SQL statement became more variable, and
 4. enclaves could result in unauthorized processor utilization.

Items 2 and 3 occurred on systems where the zIIP was a faster processor than the general CP. Item 3 also needs the application of [OA35146](#) (see earlier slide on *New Function APARs*). Most sites had rushed to install PM12256 because it improved DRDA work over TCP/IP, and it increased the portion of DRDA workloads that could run on a zIIP to 60%.

From our Tuning Letter – Great Docs

- ATS Documents (www.ibm.com/support/techdocs)
 - [WP101229](#) – *z/OS Planning Considerations for HiperDispatch Mode (Updated 13Jun2011)*
 - [WP101532](#) – *Why WebSphere Application Server for z/OS. WAS is available on platforms ranging from Windows to z/OS. Here's why you choose one over another. (Update 16Jun2011)*
 - [WP101740](#) – *WebSphere Application Server and the z/OS Workload Manager. (Updated 13Jun2011)*

From our Tuning Letter – Great Docs

- Info APAR [II07428](#) – *Reference APAR Containing Standardized Keyword Symptoms and Guidelines used When Opening or Free Form Searching APARs.*
- UK's Register had an article called "*They shoot mainframes, don't they?*" on why mainframes aren't taken seriously enough. Register to download a 9-page paper called "*Reconsider the Mainframe*" to see why consolidation of x86 servers onto a mainframe can be a win-win. The link is www.theregister.co.uk/2011/06/21/reconsidering_mainframe_whitepaper.

From our Tuning Letter – Great Docs

- IBM Redbooks (www.redbooks.ibm.com)
 - [REDP-4711-00](#) – Deploying a Cloud on IBM System z (Updated 3May2011)
 - [REDP-4727-00](#) – *Setting Up and Using the IBM System z CPU Measurement Facility with z/OS.* (11May2011)
 - [SG24-5250-08](#) – *IBM Systems Storage Solutions Handbook.* This 800-page book includes both hardware and software solutions. (26May2011)
 - [SG24-6366-02](#) – *Introduction to the New Mainframe: z/OS Basics.* (Updated 16Jun2011)
 - [SG24-7898-00](#) – *IBM z/OS V1R12 Communications Server TCP/IP Implementation: Volume 3 High Availability, Scalability, and Performance.* Provides extensive tuning recommendations. (9May2011)

z/OSMF – z/OS Management Facility

- Web-based admin tool for sysprogs
 - I've become a z/OSMF evangelist
 - Most important new facility since Workload Manager
 - IBM is firmly behind this (as evidenced by the amount of development and resources behind it) – this is NOT like eWLM, z/OS Software Manager, zMC
 - Fact – your company will need this
 - One attendee is expecting a mass exodus of MVS talent within 5 years; new sysprogs need this simplification
 - Concern – too difficult to install & too many resources to run
 - This is constantly improving, but shouldn't stop you
 - The benefits are worth the investment

z/OSMF

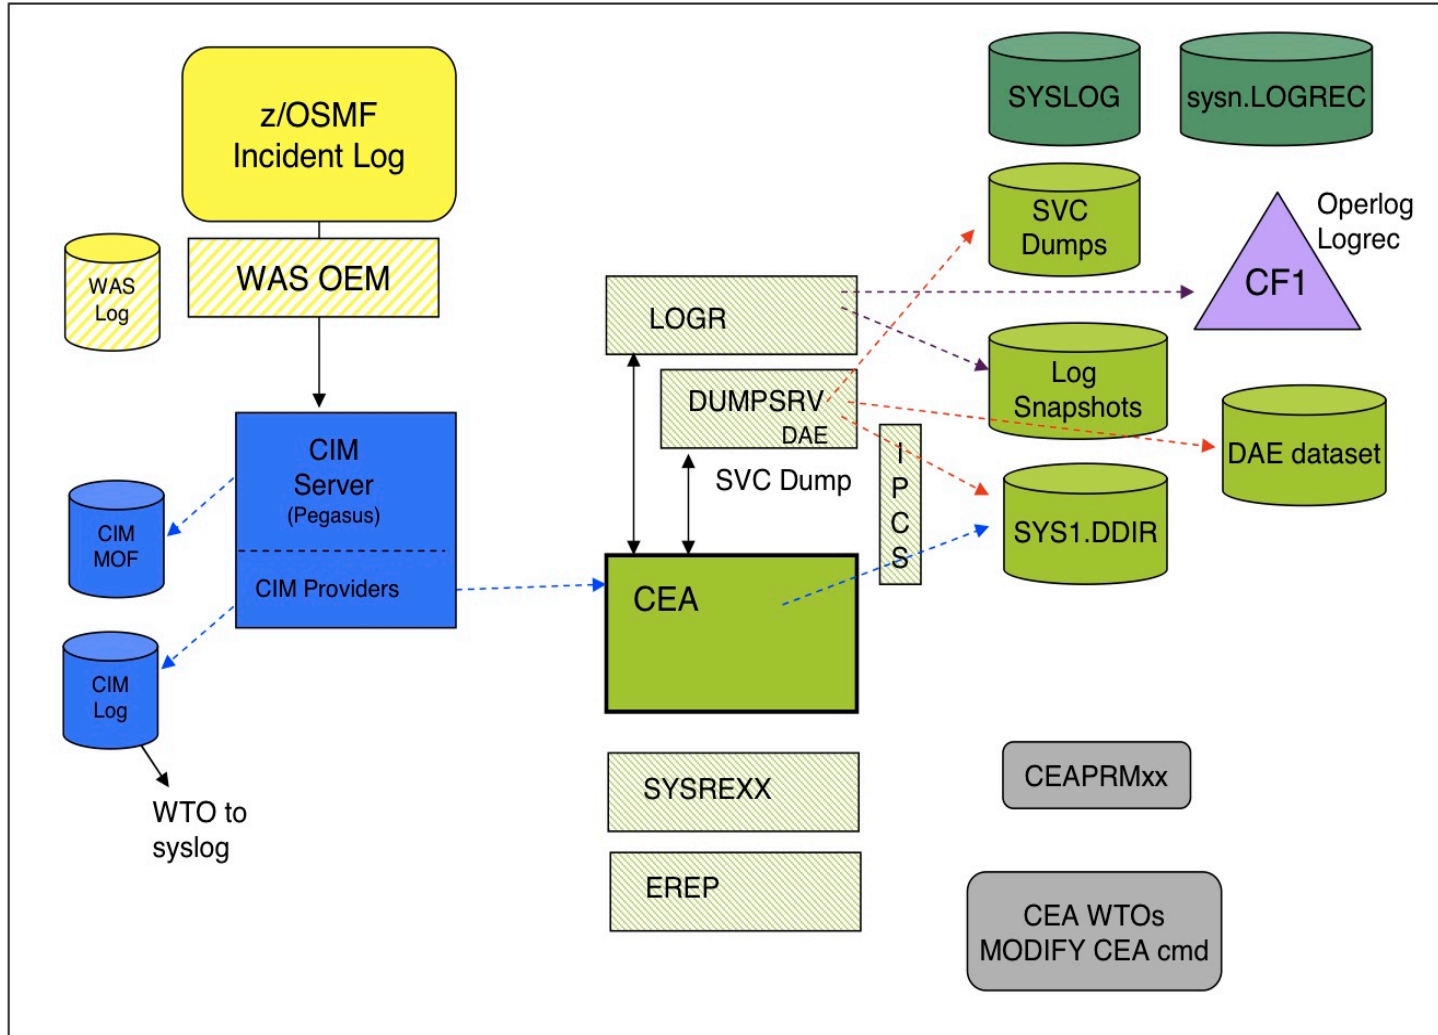
- Sessions on z/OSMF
 - 09734, [Anuja Deedwaniya](#) – *z/OSMF Overview*
 - 09804, [Anuja Deedwaniya](#) – *What's New in z/OSMF 1.13*
 - 09733, [Anuja Deedwaniya](#) – *z/OSMF 1.13 Implementation & Configuration*
 - 09737, [Mary Anne Matayaz \(Brad Carson\)](#) - *z/OSMF User Experience*
 - 10012, [Peter Muench](#) – *Manage Your Workloads & Performance with z/OSMF*
 - 09732, [Kurt Quackenbush](#) – *z/OS Software Deployment*
 - 10011, [Peter Muench](#) - *RMF: The Latest and Greatest*
 - 09735, [Anuja Deedwaniya](#) - *z/OSMF Hands-On Lab*
 - 09736 – *z/OSMF Roundtable*

z/OSMF

- Sessions on z/OSMF in Anaheim (March 2011)
 - 8696, [Karla Arndt](#) & [Anuja Deedwaniya](#) - *z/OS Problem Determination Update: z/OSMF Incident Log, Runtime Diagnostics, PFA, and New Technologies*
 - 9061, [Glenn Anderson](#) - *Understanding WebSphere App Server OEM for z/OSMF Sysprogs*
- New z/OSMF Redbook
 - [SG24-7851-00](#) – z/OS Management Facility (July 2011)
 - Also see [SG24-7853-00](#) – z/OS V1R12 Implementation (April 2011)

z/OSMF

- The following slide is taken from session 8696 earlier this year in Anaheim by [Anuja Deedwaniya](#)
 - This shows that several “best practices” should be in place before installing z/OSMF (e.g. CIM installed, CEA installed, OPERLOG operation, DAE, etc.).
 - In this slide, however, all pieces but the HMC and CIM pieces are just for use by the incident log function of z/OSMF.
 - Although I think that the incident log function is the main reason to install z/OSMF, it can also be the hardest.
 - Recommendation: Just install the WLM Policy Editor in z/OSMF 1.12 to reduce setup time. Add Incident Log later.






z/OSMF – Caution!

- [WP101779](#) – *z/OS V1.12 Management Facility Resource Requirements* (May 10, 2011)
 - z/OSMF requires a minimum of 2 GB of central storage and 120 MIPS (!!!! I gasped!)
 - But the paper also says that 3 concurrent desktops “generally require less than 1.5% of a z10 CPU”. That’s less than 14 MIPS. Why the difference? We don’t know.
 - You can offload between 25% and 75% to a zAAP.
 - Interesting: Firefox response time is better than MS Internet Explorer. You need 2 GB of memory on Windows.
 - Startup time with sufficient CPU is about 2 minutes; but can be over 20 minutes on a smaller system.

z/OSMF – Caution!

- IBM says that they will address the resource requirements, and I believe them.
- Until then, I suggest the following:
 - Understand that it will probably run like a dog (or a much slower duck) on your test LPAR!
 - Implement the easy stuff first, such as WLM and the Configuration Assistant for z/OS Comm. Server.
 - I have to believe that the “120 MIPS” is way off-base. Please send me what you think it uses at your site.
 - If using VWLC, run z/OSMF in its own LPAR so that your software costs on other software isn’t increased.



z/OSMF

- Monitor white paper [WP101779](#) for updates.
 - We've asked IBM to provide more of a breakout of resource usage by function, rather than total consumption.
- When you go to z/OS 1.13, you get so much more!
 - Software Deployment
 - Capacity Provisioning
 - DASD Management (add storage to SMS groups)
 - Browser access to ISPF
 - Enhancements to: Configuration Assistant, Incident Log, Workload Management (WLM), Resource Monitoring and System Status (RMF)

zEnterprise Issues

- z114s
- Price of specialty engines (IFLs, zIIPs, zAAPs)
- zEnterprise Positioning
- Power saving mode

z114s

- Announced in July 2011
- 130 new models to replace the z10-BCs and z9s
- Most models are positioned to match the z10-BCs, including the names. The largest PU (Z01) provides an 18% increase in speed, and the largest model (Z05) provides a 12% increase in capacity.
- Watch carefully if you move directly from a z10-BC to a z114 because there are some slight changes in MIPS and MSUs.
- Much lower cost (\$75K for the smallest A01) and lower specialty engine (\$40K for IFL, zIIP, or zAAP).

Smallest z114s

Figure 16 – z10-BC to z114 Comparison – Part 1

Model	z114 MIPS	z114 MSUs	z10 MIPS	z10 MSUs	z114 MIPS /MSU	z10 MIPS /MSU	Chg in MIPS/MSU	Chg in MIPS	% Chg in MIPS	Chg in MSU	% Chg in MSUs
A01	26	3	25	3	8.7	8.3	0.3	1	4.0	0	0.0
A02	47	6	47	6	7.8	7.8	0.0	0	0.0	0	0.0
A03	67	9	68	9	7.4	7.6	-0.1	-1	-1.5	0	0.0
A04	86	11	88	11	7.8	8.0	-0.2	-2	-2.3	0	0.0
A05	104	13	107	13	8.0	8.2	-0.2	-3	-2.8	0	0.0
B01	29	4	29	4	7.3	7.3	0.0	0	0.0	0	0.0
B02	53	7	54	7	7.6	7.7	-0.1	-1	-1.9	0	0.0
B03	75	10	78	10	7.5	7.8	-0.3	-3	-3.8	0	0.0
B04	96	12	102	13	8.0	7.8	0.2	-6	-5.9	-1	-7.7
B05	116	15	123	15	7.7	8.2	-0.5	-7	-5.7	0	0.0

Extract from Cheryl Watson's CPU Charts, July 2011

z196 Specialty Engines - IFL

Status	U.S. Price	Avg RNI MIPS (single IFL)	\$ per MIPS
New z9-BC	\$95,000	459	\$207
New z9-EC	\$125,000	560	\$223
New z10-BC	\$75,000	661	\$113
New z10-EC	\$75,000	901	\$83
New z196	\$55,000	1204	\$46
New z114	\$40,000	783	\$51
z9-EC to z10-EC Upgrade	\$0	+341	\$138
z9-EC to z196 Upgrade	\$33,000	+644	\$131
z10-EC to z196 Upgrade	\$16,500	+303	\$75

MIPS estimates come from *Cheryl Watson's CPU Charts, July 2011*

z196 Specialty Engines – zIIP/zAAP

Status	U.S. Price	Avg RNI MIPS (single PU)	\$ per MIPS
New z9-BC	\$95,000	459	\$207
New z9-EC	\$125,000	560	\$223
New z10-BC	\$125,000	661	\$189
New z10-EC	\$125,000	901	\$139
New z114	\$40,000	783	\$51
New z196	\$100,000	1204	\$83
z9-EC to z10-EC Upgrade	\$0	+341	\$139
z9-EC to z196 Upgrade	\$60,000	+644	\$153
z10-EC to z196 Upgrade	\$30,000	+303	\$128

MIPS estimates come from Cheryl Watson's CPU Charts, July 2011

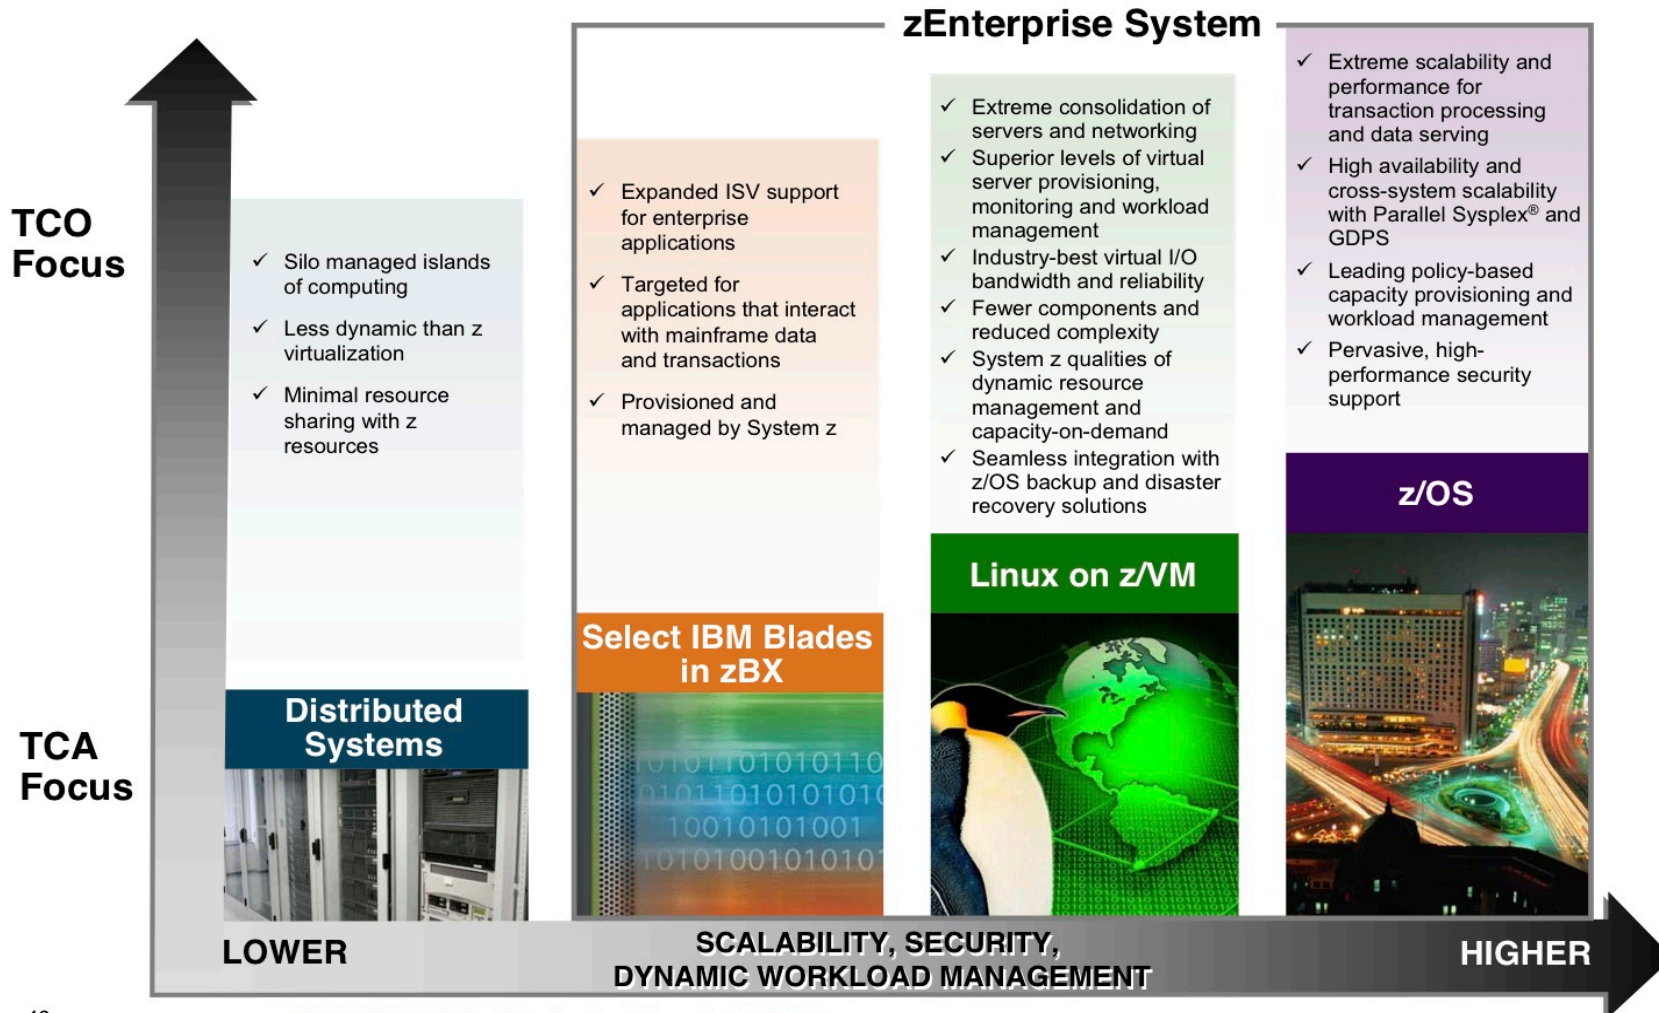
zEnterprise Positioning

- Next slide is often used by IBM to show the positioning of the zEnterprise with TCA and TCO considerations
- See Anaheim sessions 8920 by [Harv Emery](#) and 9153 by [Jose Castano](#)
- I think this is one of the most helpful displays to show the difference between Linux on z/VM and zBX



Service Levels to Match Your Business Needs

Increased flexibility for your multi-architecture strategy when data is on z/OS




z196 Power Saving Mode

- Excellent new facility on z196 to reduce power consumption when not needed (automatic or manual).
- BUT, CPU time measurements could be inaccurate.
- The service unit factor is changed, so all CPU times that are accumulated in service units are fine, but CPU times accumulated in seconds (milli-, micro-) are not modified.
- So a job that takes 10 seconds of CPU time during normal processing could take 12 seconds of CPU time during power saving mode.

z196 Power Saving Mode

- There are normalization factors in the SMF records that can be used to adjust the CPU times: 30, 70, 72, 89, 113.2, 90.34.
- See [Kathy Walsh's](#) session 09909 – *z/OS Performance Hot Topics*
- **Caution: Give chargeback and capacity planning departments plenty of notice before utilizing this feature!!**



At This SHARE

- System z Security Portal – **NEW!**
 - Security and integrity (SECINT) PTFs are included in RSUs, but cannot be found on a search of APARs.
 - You can be notified of these SECINT PTFs automatically by signing up for the System z Security Portal.
 - See www.vm.ibm.com/security/aparinfo.html for more information. You will need to get your IBM rep or IBM Business Partner to provide you approval.
 - You will not get much information other than HOLDDATA and CVSS ratings (Associated Common Vulnerability Scoring), but it's better than nothing.

At This SHARE

- 10011, [Peter Muench](#) – *RMF: The Latest & Greatest.*
 - z/OS 1.13 now collects system serialization delays (system suspend locks and GRS locks). Output is XML only, not a z/OS-viewable postprocessor report.
 - Many postprocessor reports are available via XML, eventually all will be.
 - Thanks, Peter, for the excellent notes!
- 09674, [Glenn Anderson](#) - *A Mainframe Guy Discovers Blades – as in zEnterprise “Blade” Extension*
 - Everybody needs this information.
 - Photos of wires on pages 5 & 6 brought tears to my eyes.

At This SHARE

- 09998, [Kathy Walsh](#) – *Better Batch: Exploiting New Functions to Improve Batch Processing*
 - This provides excellent suggestions on improving batch throughput.
 - As of yesterday, her handout can be found under session 09909. I expect this will be corrected soon.
 - Most people recommend that you run machines at 100% utilization, realizing that CPU times might increase. Kathy shows the much larger impact on elapsed time. She also shows what knobs to tweak to provide best throughput at 100%.

At This SHARE


- 09731, [Steve Warren, Cheryl Watson](#) – *z/OS MVSE Requirements: Influencing IBM Development*
 - The MVSE project has recently completed a cleanup and ranking of over 800 MVSE requirements.
 - This week, we presented IBM a list of our top 39 requirements (out of 209) that they will use for their planning.
 - You can help include IBM development too. See the session handout for how to join up.
- 10061/10062, [Ray Wicks](#) - *Getting Started in (z/OS) Capacity Planning (Parts 1 & 2)*
 - Ray is the master of capacity planning. Learn from the best.

From Last SHARE

- z/OS Performance Improvements
- Each new release . . .
 - Takes fewer CPU resources
 - Provides great gobs of new features
 - The following slide shows one reason you would want to keep very current with your upgrades. Slide comes from excellent session, 8861, by [Marianne Hammer](#), IBM – z/OS *Performance Issues*.

z/OS Release-to-Release Performance Improvements

- Focus on performance improvements:
 - Future z/OS releases
 - z/OS SW stack



From Last SHARE

- 9022, [John Eells](#) - *zNextGen's A SHARE'd History of the Mainframe: The Chronicles, Artifacts and Stories*
- 9028, [Mark Brooks & Nicole Fagen](#) - *Parallel Sysplex Resiliency*
 - Neat tidbit from Nicole – her favorite new IPCS command in z/OS 1.12 is 'ip systrace perfddata'. Use it for currently running system to see all address spaces and CPU usage, plus lots more
 - .
- 9066, [Glenn Anderson](#) - *A Mainframe Guy is Still Thinking About Cloud Computing*

Best Reasons to Return to SHARE

- zEnterprise
- zBX
- z196
- z/OS 1.13
- zEnterprise Unified Resource Manager)
- BladeCenter
- Blades
- Hypervisors
- Firmware
- Ensembles . . . My brain is full!

See You in Atlanta!

Cheryl Watson Walker with partner, husband, and best friend Tom Walker in Iran in May 2011
(www.tomandcheryltravels.me)



- **Email:** technical@watsonwalker.com
- **Website:** www.watsonwalker.com