



Cheryl's Hot Flashes #17

Cheryl Watson

February 16, 2007, Session 2509

Watson & Walker, Inc.

www.watsonwalker.com

home of **Cheryl Watson's TUNING Letter, CPU Chart, BoxScore, and GoalTender**

Agenda

- Survey Questions
- Rotting ROTs
- Uniprocessors
- Charge Back
- Capacity Planning
- New Function Errors
- New Latent Demand Metric
- SMF Changes & DB2 Billing
- zNextGen
- User Experiences – zIIPs
- DB2 V9
- Documentation
- Interesting APARs
- WebSphere Notes
- This SHARE

Survey Questions – Hardware

- Current Server Type (now or within next 12 months)
 - z800, z900, z890 (85% at last SHARE)?
 - z990 (50%)?
 - z9-EC (30%)?
 - z9-BC (6%)?
 - Older Hardware (4%)?
- Using zAAP Processors (25-30)?
- Using zIIP Processors (12)?
- Activated IRD CPU Management (12)?
- Have Used On/Off Capacity on Demand (12)?
- Using Variable WLC Pricing (25)?
- Doing Heavy Cryptographic Work (5)?
- Do you perform “therapeutic IPLs?”

Survey Questions – Software

- Operating System
 - z/OS.e (3)?
 - z/OS 1.4 or 1.5 (70%)?
 - z/OS 1.6 (25%)?
 - z/OS 1.7 (80%)?
 - z/OS 1.8?
 - Earlier than z/OS 1.4 (2)?
 - Note: End of Service for z/OS 1.4 & 1.5 is March 2007
- Using WebSphere on z/OS (50-60%)?
- Using VSAM RLS for CICS (10)?
- Using Transactional VSAM (0)?
- Have used zPCR (20)?
- Debug dumps with IPCS (*all but 4*)?
- Running WLM-managed initiators?
- Using RMF Monitor III?
- Using Omegamon z/OS Management Console?

Rotting ROTs

- For years performance and capacity planners have used Rules of Thumb (ROTs) for easy tuning and planning
- Are those old ROTs still valid?
- Do we still need ROTs?
- At last SHARE (and in our TUNING Letter) we proposed updating those ROTs (as part of an EWCP project)
- We'd like to retract that suggestion
- Here are the reasons why. . .

Rotting ROTs

- When ROTs first became popular, most installations had similar workloads, machine sizes, and needs
- Today's environments are extremely diverse – from 27 MIPS to over 15,000 MIPS on a single CEC
- ROTs can't be similar on such diverse configurations – the width (minimum and maximum) of each ROT would be extremely large
- Our recommendations were causing concerns in sites where the ROT didn't apply (e.g. keep CPU busy at 100%)
- So what can you use instead?

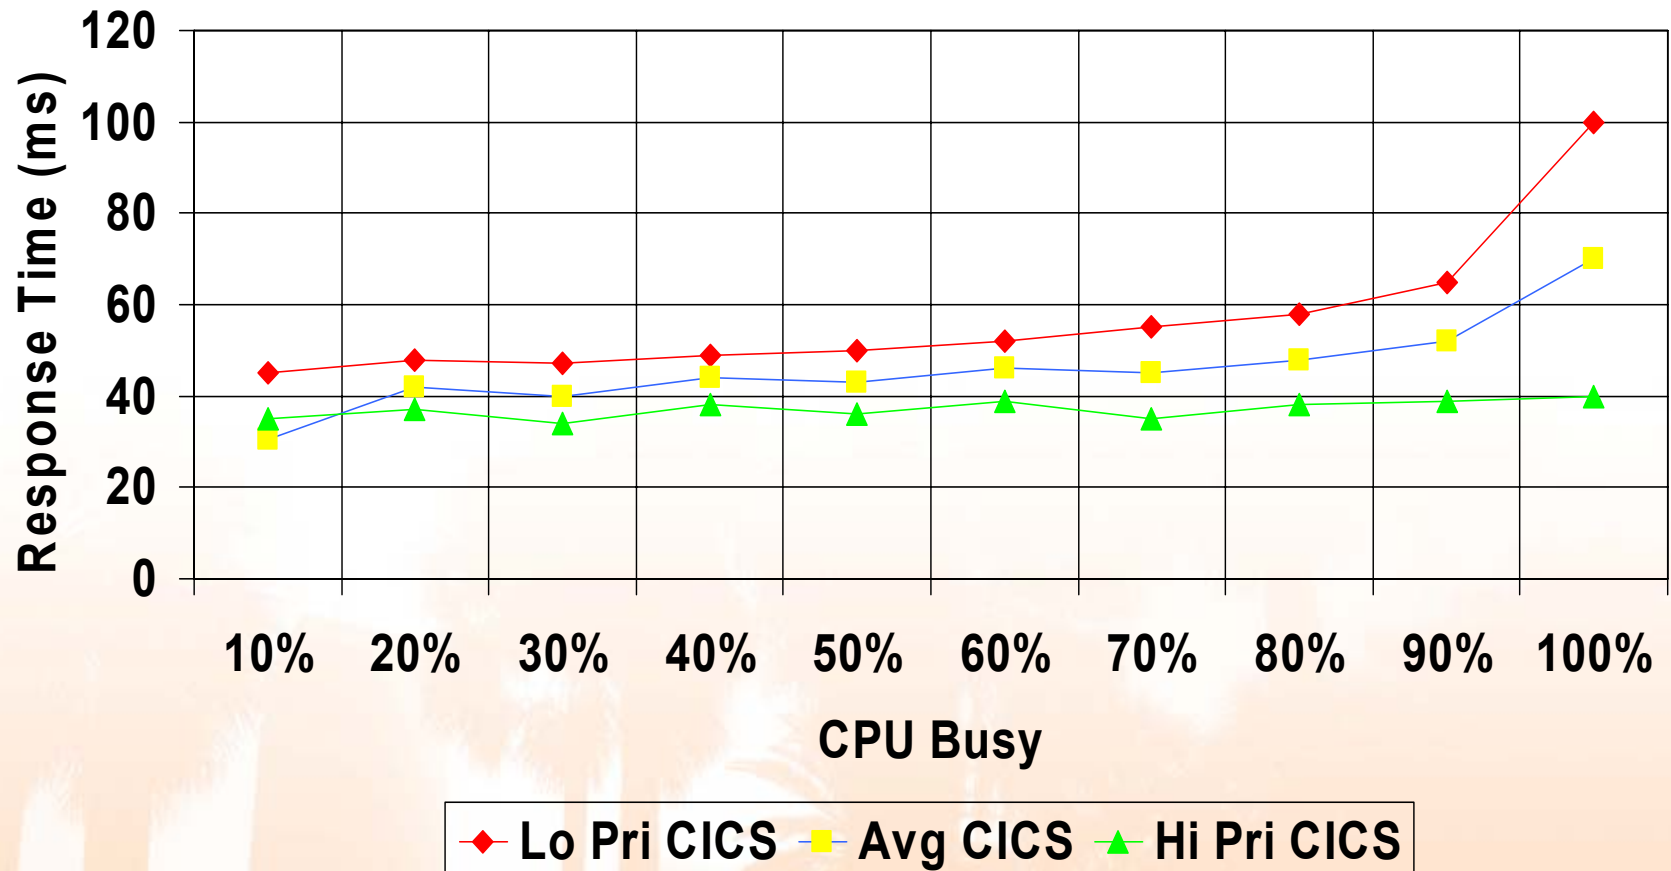
Rotting ROTs

- Two approaches can be used singly or together
- First approach – Top Ten Lists
 - 10 worst performing DASD volumes
 - 10 busiest devices
 - 10 busiest CICS transactions
 - 10 worst performing CICS transactions
 - 10 largest batch jobs for I/O activity
 - 10 largest batch jobs for CPU usage
 - Etc.
- This lets you develop good ROTs that are unique to your environment

Rotting ROTs

- Second approach – Plot metrics against response time
 - You will typically find a “knee of the curve” where the response time skyrockets when performance gets too bad
 - You need to be concerned with both high priority response times and low priority response times because Workload Manager is so good that the high priority transactions can stay responsive while the rest of the system dies
 - Sample plots:
 - Average CICS response time plotted against CPU busy
 - Better: high priority CICS transaction response time versus CPU busy
 - Even better: low priority CICS transaction response time versus CPU busy
 - Transaction response time versus average DASD service time

Rotting Rots



Uniprocessors (or small LPARs)

- Hardware Considerations
 - Initial planning for a new CEC
 - Step 1 - Estimate your needs (let's assume it's 500 MIPS)
 - Step 2 - Looking at CPU Charts, you find some possible options
 - 2094-601 – 1-way for total of 468 MIPS
 - 2094-701 – 1-way for total of 581 MIPS
 - 2096-V02 – 2-way for total of 590 MIPS
 - 2096-T02 – 2-way for total of 473 MIPS
 - 2096-R03 – 3-way for total of 554 MIPS
 - Step 3 – Run zPCR from WSC for each option
 - Step 4 – Confirm which options will provide the capacity you need
 - *Let's assume either the 2094-701 (1-way) or the 2096-R03 (3-way) meets your needs*
 - Step 5 – Review RMF/CMF data to determine appl% of applications to find whether work can use multiple CPs or will be capped by small CPs
 - Step 6 – Investigate costs of each choice (hardware, maintenance, software)
 - Step 7 – Depending on final choice, consider the items on the next foil

Uniprocessors (or small LPARs)

- Configuration Considerations
 - If 1-way CEC is chosen, carefully plan for avoiding problems with a uniprocessor
 - Two primary sources from [Linda August](#), WSC
 - *SHARE Session 2524 – “Managing Workloads on a Uniprocessor”*
 - *WSC White Paper - **WP100925** – “Managing CPU-Intensive Work on Uniprocessor LPARs” (15Dec2006)*
 - If 2-way CEC is chosen, you have two options for small LPARs:
 - Avoid uniprocessor problems by assigning a minimum of two LPs to each LPAR – but the cost of this solution is CPU overhead
 - Assign a single LP and cover your bases with the same procedure as above (review Linda’s work)
 - Maybe best of both worlds: assign two LPs at IPL, followed immediately by a vary CPU offline; this allows the operator to vary another LP online in the case of a loop

Charge Back

- How many sites do charge back to internal users?
- How many sites do charge back to external users?
- Two options currently used:
 - CPU Time – great loss of time due to lack of precision in type 30 records
 - Service Units Consumed – varies by 20 to 40% depending on many things: number of LPs assigned to LPAR; number of CPs (of all types) on CEC; utilization of the system (i.e. CPU busy); type of other work on the system; time of day, etc.
- Recommendation:
 - Use “bands” of resource usage for charge back (e.g. 0 to 10 seconds CPU time = \$xx; 11 to 30 seconds = \$yy; 31 to 60 seconds = \$zz)

Capacity Planning

- Majority of capacity plans use service units
- Service units aren't consistent
 - Small LPARs on a large CEC appear to take less time (based on the SU/Sec value assigned), but in fact take more CPU time due to overhead from other LPARs
 - Service unit usage increases as the CEC utilization increases, especially for memory hogs
 - Service units per second are constant when using IRD, even as LPs are varied on and offline
 - Variation used to be within 5%, now it's closer to 20-30%
 - What to do?
 - EWCP project is asking IBM for help and direction
 - We are submitting a requirement that all LPARs on a CEC use the CEC value of SU/Sec instead of the LPAR view of SU/Sec

New Function Errors

- IBM policy (since early 2000s) causes new function maintenance to bypass PE status, even when known errors exist.
- Discovered by a customer when researching whether to apply the **OA12364**/UA90255 fix to allow use of JES2 NJE over TCP/IP with z/OS V1R7. There was already one PTF (UA28508) and one open APAR (**OA17300**) caused by UA90255. The original was not marked as PE.

Thanks to **Michaela Weber**, office of the Idaho State Controller

New Function Errors

- IBM's Response:
 - Even if the new APAR is an error, the new function APAR may not be marked PE if it follows these guidelines as officially documented in the PE PDR:
 - If the error is injected by a new function APAR (SPE) and the error is neither HIPER, nor can it be discovered by a customer without taking overt action to use the new function (e.g., the error does not regress existing function), then the error will not be considered a PE.
 - One of the major reasons for not marking a NF PTF as PE is that marking it PE could prevent customers from installing corrective service or other critical service for an error that you can't hit, because you haven't enabled the function.
- “Overt Action” is more than just applying the APAR
 - For example, coding a new PARMLIB option to activate the feature
- We understand IBM's logic, but you need to be aware of this

New Latent Demand Metric

- 'In and Ready' greater than the number of LPs shows the number of ready tasks that aren't getting dispatched
- For years this field has been tracked to indicate latent demand
- But the number of LPs now changes dynamically due to: IRD dynamic CPU management, operator vary commands, CPU on demand, etc.
- z/OS 1.7 provides new data fields in type 70 record:
 - SMF70Q00 to SMF70Q12 – Count of the In Ready users based on the number N of CPs being online when the sample was taken
 - SMF70Q00 – Count of users less than or equal to N
 - SMF70Q01 – Count of users equal to N + 1
 - SMF70Q02 – Count of users equal to N + 2
 - SMF70Q03 – Count of users equal to N + 3
 - SMF70Q04 – Count of users equal to N + 4 or N + 5
 - SMF70Q05 – Count of users equal to N + 6 to N + 10
 - ...
 - SMF70Q12 – Count of users greater than N + 80
- New metric for latent demand = $((SMF70Q01 * 1) + (SMF70Q02 * 2) + (SMF70Q03 * 3) + (SMF70Q04 * 4.5) + (SMF70Q05 * 7) \dots (SMF70Q12 * 80)) / \text{sum}(SMF70Q00 \dots SMF70Q12)$

SMF Changes and DB2 Billing

- If you use the Usage License Charge (ULC) method of billing you may see pricing changes after applying a new APAR, particularly if you run DB2.
- **OA17891** (z/OS 1.4+, 7Nov2006) - *Corrections Are Made to the Accumulation of Product Usage-Related Fields within the SMF89 and SMF30 Records.* Product usage data that is recorded in the SMF type 30 and type 89 records may contain incorrect values. These values can be found in the fields SMF89UCT (product TCB time), SMF89USR (product SRB time), SMF30UCT (product TCB time), and SMF30UCS (product SRB time).
- The fields now include CPU time for enclaves (task or SRB), client SRBs and other pre-emptible SRBs, so can be much higher than previously reported.

Thanks to Brian Peterson from Blue Cross Blue Shield of Minnesota

zNextGen

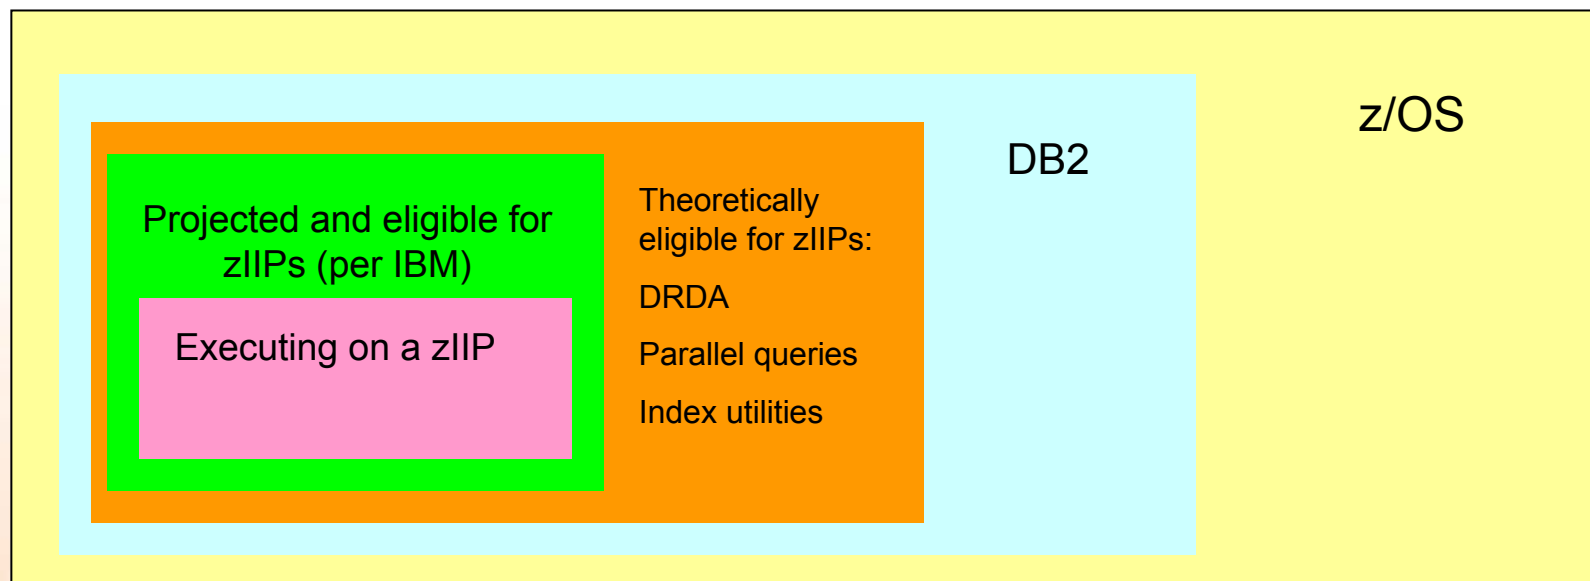


- To reverse the graying of SHARE and our industry, please promote zNextGen in your company
- This is a new SHARE project that can help train young people working with the mainframe and more mature people who are transferring job skills
- Many initiatives, including university training, online training, help sites, etc.
- Best thing you can do for your replacement is to get them to the next SHARE

- z/OS Basic Skills Information Center - publib.boulder.ibm.com/infocenter/zoslnctr/v1r7/index.jsp
 - **z/OS concepts**
 - **Application programming on z/OS**
 - **Networking on z/OS**
 - **Problem management on z/OS**
 - **Security on z/OS**
 - **System programming on z/OS**
 - **Online workloads for z/OS**
 - **Interactive courses**
 - **Glossary of z/OS terms and abbreviations**
 - **Help for the information center**
 - **Contact z/OS**

User Experiences - zIIPs

- Word of mouth: zIIPs are outperforming expectations of users; customers are very happy with results
- Definitions on chart – each box is subset of larger box
 - We have all measurements but ‘Theoretically eligible’



User Experiences - zIIPs

- How much can your installation move to zIIPs?
- Only benefits now are seen in large DB2 sites with DRDA (distributed DB2), utilities, and large parallel queries
- Future applications on zIIPs: TCP/IP Sec (security) and more
- Important to run projection tools first and to look at RMF data
- The pricing benefits of zIIPs depends on what's causing the highest CPU usage by time (rolling 4-hour average)
 - If it's DB2 that could run on a zIIP, then a zIIP can reduce your cost
 - If there is no 'eligible' DB2 during your peak hours, then a zIIP may not help out at all
- Requirements: DB2 V8, System z9, z/OS 1.6, zIIP FMID

For more information on zIIP measurements, see [Kathy Walsh's](#) session 2512; [Walt Caprice's](#) session 2515; and [Peter Enrico's](#) session 2523

DB2 V9

- See session 1300 – “What’s New in DB2 for z/OS V9 and Beyond”
- Key points to know:
 - DB2 V7 has end of service June 2008
 - DB2 V8 had increase (5-10%) in resource usage, which can be eliminated by exploiting many of the new functions
 - DB2 V9 has decrease in resource usage:
 - Up to 30% reduction in many utilities
 - Improved index compression - one site saw 50% reduction
 - z800 and z900 could see 5-10% degradation instead of improvement
 - V8 SQL procedures couldn’t use zIIP, but native SQL Procedure Language on V9 can use zIIP
 - When will V9 be ready? Beta began in June 2006; wait & see; but V9 migration only available from V8 NFM

Neat Documentation

- **SG24-7324-0** - IMS Performance and Tuning Guide (18Dec2006)
- **SG24-7225-0** - z/OS V1R7 DFSMS Technical Update (29Dec2006)
- **Tips0647** - Migrating from Hierarchical File Systems to zSeries File Systems (24Jan2007)
- **TD103468** - z/OS 1.8 Installation Checklist (30Nov2006)

Interesting APARs

- **OA19257** (z/OS 1.6+, 10Jan2007) - *CPU and SRB Service Units Inaccurate Due to Truncation*. Due to truncation, You may see inaccurate and smaller (even negative) accumulated times in SMF30CSU, SMF30SRB, R723CCPU, R723CSRB, RQSVCPU, RQSVSRB, RCAECPU, and RCAESRB.
- **OA12835** (z/OS 1.5-1.7, 19Aug2005) - *High CPU Usage, MQ CHIN Task, z/OS R1.5-1.7, IGVFSDQE SYSTRACE EXT and CLKC Events During FREEMAINs SP000/Key8*. Although this HIPER maintenance is not new, Brian suggested that many users may have overlooked it, because they thought it only applied to WebSphere MQ. But the performance benefit should apply to any application that is a large consumer of storage.
- **OA19462** (z/OS 1.6+, 23Jan2007) - *SRB Service Units Incorrect at JBB77S9, HBB772S and HBB7730*. In z/OS 1.8 and 1.6/1.7 with the FMID for zIIPs applied, the consumed SRB service units are multiplied by the CPU coefficient instead of the SRB coefficient.

Thanks to **Brian Currah** of BDC Computer Services

Interesting APARs

- CATALOG (12Feb2007)
 - Opened just last Monday!
 - A sysplex-wide hang can occur waiting for catalog resource SYSIGGV2 after application of **OA18860** (16Dec2006). Affects those running z/OS 1.4 - z/OS 1.8.
- WLM DP (27Dec2006)
 - **OA18418** - DISPATCHING PRIORITY OF SERVICE CLASSES ARE LOWERED. Performance problems could occur because the dispatching priorities could be the same for any service class regardless of their importance. Affects z/OS 1.6 – z/OS 1.8.
- DB2 V8 High CPU (9Dec2006)
 - **PK33403** - HIGH CPU IN CUNMUNI WHEN RUNNING SQL HEAVY APPLICATION UNDER DB2 V8. Possible performance problems using the CICS/DB2 attachment with DB2 V8 running NEWFUN(YES). CPU cost per SQL of a program pre-compiled using DB2 V8's NEWFUN (New Function) mode increases at CICS/TS 3.1 following APAR **PK05933**.

Thanks to **Jerry Urbaniak** of Acxiom

WebSphere Note

- We hear from many installations who want to implement WebSphere on z/OS
- Most don't have the time, and certainly don't have the knowledge or experience
- One easy entry into Web serving is to use HATS (Host Access Transformation Services) - described in our Hot Flashes #15. You could have a Web application converted from green screens within a week. (Also see session 3415 by [Anthony Lewitt](#) this SHARE.)
- For actual WebSphere solution to allow development of new applications: contact your IBM rep for help
 - If you can justify your need and interest in WebSphere, IBM might provide free onsite implementation help (this also takes about a week)

This SHARE

- **Interesting Sessions**

- 2500, “zOS Performance ‘Hot’ Topics” by **Kathy Walsh** – latest in WSC information and performance APARs (always my favorite session at SHARE).
- 2840, “Building a ‘Gee Whiz’ z//OS Toolkit” by **Clark Kidd**, Watson & Walker, Inc. – see www.watsonwalker.com/toolkit.html
- 3027, “Look What You Can Do with DFSORT and ICETOOL Now!” by **R. David Boenig**
- 1314, “A DB2 Performance Tuning Roadmap” by **Craig S. Mullins**, Neon Software
- 2515, “WSC Experiences with the zIIP Processor” by **Walt Caprice**
- 2552, “CPU reporting in z/OS 1.7 RMF and SMF” by **Linda August**
- 2530, “z/OS Performance Update V1.7 and V1.8” by **Marianne Hammer**

This SHARE

- RMF Monitor III Data Portal
 - Gives you Web browser view to RMF Monitor III data
 - The facility is free (comes as part of RMF)
 - Available on z/OS 1.8, 1.7 with UA90253, and 1.51.6 with UA90251
 - Very easy to implement
 - MUCH easier to view and use than RMF MIII
 - Additions:
 - Graphical view of any of the data elements
 - Easy sorting of columns, plus scrolling!
 - Customizable view to show multiple key elements
 - References:
 - Hot Topics, pages 56-57
 - SHARE Session 2557 by **Peter Muench**
 - RMF Website: www.ibm.com/servers/eserver/zseries/zos/rmf/

This SHARE

- IBM Migration Checker
 - Free downloadable utility
 - Similar (in concept) to the Health Checker
 - Checks the status of z/OS migration tasks
 - Provides recommendations
 - Makes NO changes to your system
 - Designed for z/OS 1.7 to z/OS 1.8 migration
 - But some checks work with earlier z/OS releases
 - There are 12 checks in the first version
 - For more information...
 - Session 2870; Migrating to z/OS R8 Part 1 of 3; **Marna Walle**
 - Session 2800; MVS/SCP Project Opening & Hot Topics; **Bette Brody**
 - Hot Topics; February 2007, pages 14-15
 - Our next TUNING Letter

This SHARE

- Preview of z/OS 1.9 (What we like)
 - Constraint relief in GRS (more control blocks above the bar)
 - TSO/E support for large sequential data sets (> 64K tracks)
 - Support for SMF records written to Logger
 - More flexibility
 - Better performance (4x to 7x faster in some IBM tests)
 - Message flood automation (control runaway console messages)
 - Reduced z/OS UNIX latch contention
 - Better detection/recovery of unresponsive systems in a sysplex
 - Workload Manager enhancements
 - Improved routing of zAAP and zIIP workloads
 - Addition of “trickle” support for low-priority work
 - Promotion of cancelled tasks
 - UNIX file/directory deletion recorded in SMF 92 record

This SHARE

- Preview of z/OS 1.9 (What we like...continued)
 - System REXX
 - Execution of REXX routines in an authorized environment
 - Operator execution of REXX utility routines
 - Future support for Health Checks written in REXX
 - More Health Checks
 - ISPF edit and browse of UNIX files
 - RMF reporting of coupling facility utilization
 - Support of masking characters in IDCAMS DELETE
 - Improvements to TRSMAIN
 - Support for large format sequential data sets
 - Other improvements
 - Will become part of z/OS, and will be supported
 - Binder validates RECFM=U on output data set

See You in San Diego!



- **Email:** technical@watsonwalker.com
- **Web site:** www.watsonwalker.com