

Cheryl's Hot Flashes #12

Cheryl Watson

Session 2509; SHARE 103 in New York City

August 20, 2004

Watson & Walker, Inc.

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

Ain't it the Truth!



- And you thought that Y2K was a killer:
 - *"IBM has another timing problem of even greater impact in the somewhat distant future sometime in 2042 when the current 8-byte TOD clock wraps."*
(Jerry Urbaniak)

- Taglines
 - *"One of the main causes of the fall of the Roman Empire was that, lacking zero, they had no way to indicate successful termination of their C programs"*
- *Robert Firth* (Murray Butler)

Agenda



- Survey Questions
- Interesting Items – Neat Info
 - System symbols, WebSphere, Publications, CICS, Under the Covers, APARs
- User Experiences
 - Striped Data Sets, WLM Overhead, 9672s, Chargeback
- 6-Month Update
 - 31-bit on zSeries, SMB, SMF, z990 Sizing
- This SHARE
 - Neat Info, zAAP Processors

Survey Questions



- Who is running z/OS.e?
- Who is using Variable WLC pricing?
- Who has activated IRD CPU management?
- Is anybody still running OS/390?
 - If so, be sure to order z/OS 1.4 or 1.5 before Sep2004
- Is anybody still running G5/G6 machines?
 - If so, be sure to order z/OS 1.4 or 1.5 before Sep2004

Survey Questions



- Who is planning to investigate On/Off Capacity on Demand in the next year?
- Who is looking into zAAP processors?
- Who is running WebSphere Application Server on z/OS?
- Who uses SMF type 30 CPU time for chargeback, data center reporting, pseudo-chargeback, capacity planning, etc.?
- Who is doing heavy cryptographic work?

Interesting Items – Neat Info



➤ System Symbols

- Can you dynamically create your own system symbols and dynamically change existing symbols?
- Official answer is no
- But an old Redbook provides a utility to perform this function (has been run on z/OS 1.4) - (**SG24-5451-00, Parallel Sysplex - Managing Software for Availability**, Dec1999)
- Link in Redbook provides download for free utility

Interesting Items – Neat Info



➤ WSC TechDocs

– Lots of information on WebSphere – Hints & Tips

- **TD100434** - Finding and Collecting Availability Measurement Data for the IBM HTTP Server and WebSphere Application Server on z/OS (13Apr2004)
- **TD100745** - Activating S/390 and zSeries Cryptographic Services for WebSphere (25Mar2004)
- **TD101118** - RACF Tips for customizing WebSphere for z/OS Version 5.0.2 (16Mar2004)
- **TD101151** - How to Classify HTTP Transactions in WebSphere for z/OS V5 (1Jun2004)
- **TD101246** - Log4j Enhancements for J2EE Applications under WebSphere Application Server v5 for z/OS (25Mar2004)
- **TD101255** - Implementing Enhanced Form Based Authentication with Servlet Filters in J2EE Applications under WebSphere Application Server for z/OS (18May2004)
- **TD101529** - Application Migration Perform Guide - Migrating to WebSphere Application Server 5.0.2 for z/OS (1Apr2004)
- **TD101663** - Enabling WebSphere Application Server 5.0.2 for z/OS to use the DB2 Universal JDBC Driver (28Apr2004)
- **TD101703** - Disabling the Deployment Manager HTTP Timeout (18May2004)
- **TD101711** - Integrating WebSphere Applications with IBM Tivoli Enterprise Console or IBM NetView for z/OS for Event Management (25May2004)

Interesting Items – Neat Info



➤ WSC TechDocs

- Lots of information on WebSphere – White Papers
 - **WP100392** - Exploiting Web services in WebSphere for z/OS (7May2004)
 - **WP100415** - WebSphere V5 for z/OS -- Starting Deployment Manager on Another MVS Image (6Apr2004)
 - **WP100421** - WebSphere Version 5 for z/OS - WSADMIN Primer (5May2004)
 - **WP100424** - WebSphere Application Server V5 for z/OS JMS and MDB IVP (11May2004)
 - **WP100426** - WebSphere Portal with Tivoli Access Manager: Value Beyond Security (25May2004)
 - **WP100431** - Installation of the zAAP Projection Tool Instrumented SDK in WebSphere for z/OS Version 5 (2Jun2004)

Interesting Items – Neat Info



➤ Publications

- One of people's favorite tips from last SHARE
- Great publications site: www.elink.ibm.link.ibm.com
- This is IBM's site for EMEA (Europe, Middle East, Africa)
- For the USA, go directly to www.elink.ibm.link.ibm.com/public/applications/publications/cgi-bin/pbi.cgi?CTY=US
- It's a lot easier to use than IBM's regular library site at www.ibm.com/servers/eserver/zseries/zos/bkserv

Interesting Items – Neat Info



➤ CICS Document

- **Bob Yelavich** publishes his CICS newsletter, *Follow the Yelavich Road*, on his Web site at www.yelavich.com
- The 06Jan2004 issue had an extremely interesting article on CICS and multiprocessor exploitation. If you've been frustrated that CICS could seldom use more than one CP, then this article is for you. CICS TS 2.2 has the capability to exploit multiple CPs if you're willing to give it a try. For a direct link, use www.yelavich.com/cn010604.htm.

Interesting Items – Neat Info



➤ Under the Covers

- Great section on how LPAR dispatching works:
 - **SG24-5952 - z/OS Intelligent Resource Director**
- Great description of how dispatcher works, starting in MVS SP 5:
 - **GG24-4137 - MVS/ESA 5.1.0 Technical Presentation Guide**
 - *'Dispatching Management in MVS: TCBs to Enclaves'*, **Bernard R. Pierce**, CMG Transactions, August 1996
 - *'The Evolution of the SRM to The Workload Manager in MVS V5'*, **Bernard R. Pierce**, CMG Transactions, Winter 1995
 - Also: **Cheryl Watson's** TUNING Letter, 1999, No. 2

➤ OA06611 – GRS/MIM

- Open (16Feb2004) APAR to improve GRS/MIM performance
- Reduces pathlength by eliminating some of the prebatch exit calls
- CA-MIM 4.6 supports this facility (need PTF for high CPU time)
- Wells Fargo implemented parts of new APAR, CA-MIM updates and application changes to remove enqueue/dequeue to a dataspace – results are that CA-MIM/GRS performs better in z/OS 1.4 than in OS/390 R10

➤ WLM Managed Initiators

– OA06350

- WLM Managed Initiators Not Starting
- OS/390 R10-z/OS 1.5, 17Mar2004

– OA06687

- IWMRCOLL RCAEQ AND RCAEUNKN FIELDS MAY CONTAIN INCORRECT VALUES
- OS/390 R10-z/OS 1.5, 28Apr2004

– Thanks to **Patrick Falcone** of Combined Insurance

➤ GRSQ Performance Enhancement

– OA07975

- OPEN, 10Jun2004
- CICS dumps were taking 10-30 minutes and using over 10 million service units
- Usermod to bypass collection of GRSQ information dropped elapsed time to 1-2 minutes and using 90% fewer service units

Thanks to **Curt Thompson** of the Principal Financial Group

- VLF/RACF Performance Improvement
 - APAR [OA06920](#) (OS/390 R10+, 17Jul2004)
 - In some situations, a VLF timing issue may prevent it from caching certain RACF data structures. This can cause increased CPU usage and performance problems with any application that is waiting on RACF. Although the problem can affect any VLF user, it seems to be RACF that is particularly vulnerable.

Thanks to **Tom Conley** of Pinnacle Consulting Group for bringing this to our attention.

User Experiences – Striped Data Sets



➤ z/OS 1.5

- Non-VSAM extended data sets can now reside on 59 volumes instead of 16 volumes
- **Steve Mawhinney** of Wells Fargo shared their experiences under OS/390 R10 and z/OS 1.4 with a usermod to provide this same function:
 - 1.3 million track data set with 310 jobs accessing it a day – the average run-time is 92% less when using 48 stripes
- Be sure to check out this function for large, multi-access data sets!

User Experiences – WLM Overhead



➤ Enclave Explosion

- This was discussed by **Laura Blodgett** of IBM (session 2544)
- ENVT control block is created for each enclave in the system and sampled by WLM
- Size of sampling storage area is based on the number of ENVTs and ASVTs; new storage area is created when more enclaves are created; space is never released, so grows from start of IPL; clearing of new storage is done with interruptible MVCL; can cause overhead in WLM address space; Laura says this can be as high as 10%
- Word of mouth: A user in the same session said that ***they lost half of a z990 engine in a 6-way machine*** due to this problem
- IPL reduces the overhead
- Also see WebSphere APAR [PQ84981](#)

User Experiences – WLM Overhead



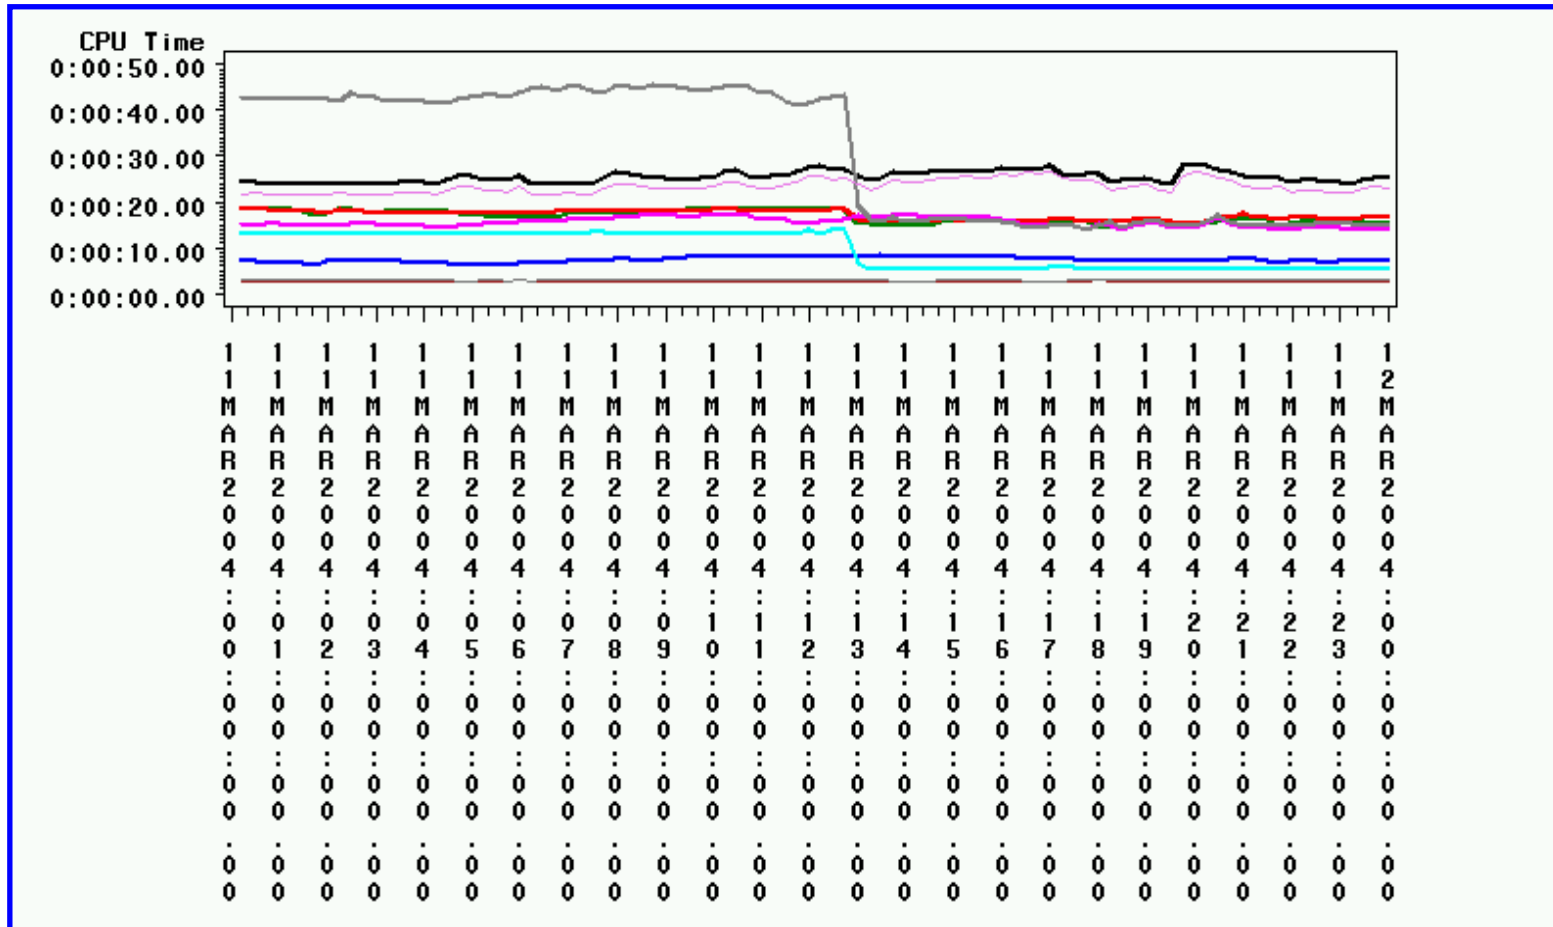
➤ CICS PBs

- If you manage a CICS region with transaction goals, there is one Performance Block (PB) built for each MAXTASK value
- WLM then samples these PBs every quarter second
- Recommendation – set 'Region' control for any region that does not need to be managed by transaction goals (such as test regions). Then manage them by velocity. This is usually done in the STC classification rules.
- Example: **Chuck Hopf** saved 35 seconds (from the WLM address space) every 15 minutes by switching to region control in selected regions (see next foil for results)
- Use similar plot to see if CPU time is increasing during life of IPL

User Experiences – WLM Overhead



Figure - WLM CPU Time



User Experiences – 9672s



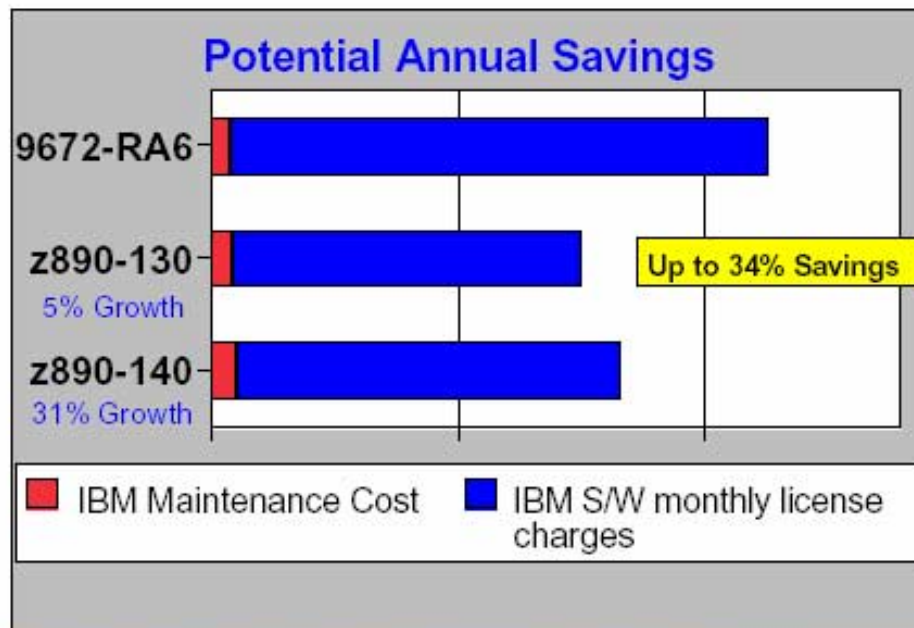
- z/OS 1.6 *ONLY* runs on zSeries machines; z/OS 1.4 EOS is Mar2007
- Have you looked at cost of z890?
- People tell us that the z890 pricing with its variety of MSUs can be cheaper than the 9672 (due to hardware maintenance and software pricing)
- At least give it a look – it's crazy not to!
- The following slide is courtesy of **Mike Augustine** of IBM

Replace your 9672 G5 running OS/390 to a z890 with z/OS and Save!



▪ Innovation on z890

- ƒ Latest 64-bit architecture
- ƒ zAAPs: zSeries Application Assist Processors for cost-effective, high performance Java applications next to Legacy data
- ƒ High Capacity IFL
- ƒ Improved I/O Connectivity with OSA-Express, OSA-ICC, HiperSockets
- ƒ On/Off Capacity on Demand
- ƒ PCI Crypto



▪ **Ongoing Savings help you offset the cost of your hardware investment.**

▪ Savings driven by:

- ƒ Lower priced z/OS on the 130 or 140 with full capacity EWLC versus OS/390 with PSLC software pricing
- ƒ Lower software pricing on certain other IBM MLC software for z890

▪ Additional Value considerations:

- ƒ Granular upgradeability
- ƒ Upgradeable to z990

* Hardware configuration based on general purpose central processors only (no IFLs, zAAPs or CFs), and similarly configured for costing purposes. z/OS software stack based on z/OS, including DFDSS, DFHSM, DFSORT, RMF, SDSF and Security Server, CICS, DITTO and UDB.

User Experience - Chargeback



➤ Chargeback Issues

- With variable z990 CPU times and zAAP processors, repeatable (or even fairly consistent) chargeback costs are becoming more difficult to obtain
- A change in the amount of storage can also dramatically change the CPU time usage of a job
- Consider tier/range charging:
 - CPU time between 0-10 seconds = \$1.50
 - CPU time between 10-60 seconds = \$5.00
 - CPU time between 60-300 seconds = \$10.00
 - Etc., etc.
- Variability then won't impact the customers as often

6-Month Update – 31-bit on zSeries



➤ 31-Bit on zSeries

– Common scenario

- OS/390 R10 running on non-zSeries and upgraded to zSeries
- Normally storage increase is very large

– Example

- Site moved from 9672 with 1 GB of storage to z800 with 8 GBs of storage (minimum configuration)
- CPU time in *MASTER* increased from 1% to 10%
- Cause is due to additional scanning of expanded and frequency of central scanning (every 2 seconds)

6-Month Update – 31-bit on zSeries



➤ 31-Bit on zSeries

– Solution

- Vary storage offline until it's needed
- Migrate to 64-bit using bimodal mode
 - 64-bit will eliminate scanning of expanded and frequency of central storage scanning decreases to every 10 seconds
- Use SoftCap from WSC ([PRS268](#)) to estimate cost of move

6-Month Update – zSeries



➤ Speaking of Larger Storage

– RSU=xx

- Parameter in IEASYSxx
- Used to be in storage increments
- In OS/390 R10, you can also specify MB, GB, %
- **Please, please, please stop using increments!**
- RSU=2 means:
 - 32 MB on smallest z800
 - 128 MB on z900 with 32 GB
 - 256 MB on z990 with 40 GB

6-Month Update – SMB



➤ System Managed Buffering (SMB)

– APAR [OW41269](#)

- Documentation For System Managed Buffering (SMB), 09Nov1999
- Excellent guidelines and restrictions for SMB
- Incorrect last SHARE: "Was supposed to be added to **Using Data Sets** manual, but hasn't made it yet"

Thanks to **Mark Thomen** of IBM for pointing that this information **is** contained in the manual

6-Month Update – SMF



- APAR [OA06407](#) – Large or negative values in SMF30SRV
 - Closed 18Feb2004 for OS/390 R10+
 - **Kathy Walsh** mentioned this in session 2500
 - What we saw on z900:
 - The 16 largest job steps out of 1,069,999 had CPU time greater than the service unit time (404 hours CPU time versus 119 hours SU time)
 - So before using service units, be sure to apply the fix for this APAR!

6-Month Update – z990 Sizing



- SHARE Session 2513 – Introduction to Processor Sizing, **Joanne Brown**, WSC
 - New LSPR workloads for zSeries (primitives):
 - CB-L – Commercial Batch (similar to CBW2)
 - CB-S – Commercial Batch Short (similar to CB84)
 - OLTP-W – Online (similar to CICS/DB2)
 - OLTP-T – Online (similar to IMS)
 - WASDB – WebSphere
 - Important to use for sizing z990s

6-Month Update – z990 Sizing



➤ From last SHARE -

- Marketing reps can use free internal tools for processor sizing studies (zPCR and CP2000)
- Three new custom workloads:
 - Online: $(\text{OLTP-W} + \text{OLTP-T} + \text{WASDB}) / 3$
 - Other: $(\text{CB-L} + \text{CB-S}) / 2$
 - Low I/O: $(.60 * \text{CB-L}) + (.20 * \text{OLTP-W}) + (.20 * \text{WASDB})$
 - Low I/O is defined as images that have fewer than 30 DASD SSCHs/Second per used MSU
- Most installations are in a low I/O environment

6-Month Update – z990 Sizing



- This SHARE – My session 2537 – Are z990s Underperforming?
 - If you don't use zPCR or CP2000, you may be disappointed
 - IBM rep is required to run one of these tools prior to completing a contract for a z990
 - Our experience – the people at the site who confirm the capacity of a new machine have never heard of these tools
 - **It's imperative that you get the tools and understand the results before confirming your hardware order**

6-Month Update – z990 Sizing



- This SHARE – My session 2537 – Are z990s Underperforming?
 - I chided IBM for their +/- 5% variance on the zPCR results because -5% can amount to a few hundred MIPS!
 - **Kathy Walsh** of WSC later pointed out that the tools really can't be more accurate than that due to the many variables involved (as I mentioned in my presentation): 64-bit, fewer but faster CPs, uni-processor considerations, higher LP to CP ratios, difference in LPAR configurations, etc.
 - We now understand why the +/-5% variation, but what can you do?

6-Month Update – z990 Sizing



➤ z990 Sizing Considerations

- From our analysis, low I/O sites going from a z900 to a z990 are typically seeing about 10% underperformance if using MIX MIPS (using low I/O customized workload, then it's fairly close)
- From our experiences, the z990s produce the greatest variation of any processor we've ever seen
- Based on IBM's zPCR sizing, you should be sure you can handle a capacity that is 5% less than you expect, even when using low I/O
- Are z990s Underperforming? – The answer is NO, as long as you size correctly, tune your system, adjust your LPAR configurations and consider storage changes.

6-Month Update – z990 Sizing



- Gartner publishes single number MIPS value
 - This week they dropped their published MIPS by about 10%, which is probably to account for the low I/O effect (10% apparent underperformance if using MIX or average MIPS when moving from a z900 to z990)
 - We disagree with a single published MIPS value; especially now that customized workloads are more applicable than MIX MIPS
 - We think, instead, that ISVs should be using the published software MSUs for charging instead of MIPS (these happen to be 10% below hardware MSUs)
 - Because of low I/O, you may not see a drop in software costs on the z990s

This SHARE – Neat Info



- Informative Session – 2544
 - **Laura Blodgett** of IBM, "WLM Bits and Pieces"
 - Presentation describes WLM's logic for policy adjustment and resource group analysis
 - Good APAR –
 - [OA03207](#) (OS/390 R10 – z/OS 1.5, 25Mar2003) – keeps 15 minutes worth of type 99 data internally despite whether 99s are written to SMF (useful for debugging)

This SHARE – Neat Info



- Storage Administration z/OS Pocket Reference
 - DTS Software publishes this
 - Get free one from www.DTSsoftware.com
- The Case For and Value of Transactional VSAM
 - **Terri Menendez** – Session 3064 (see also 3063)
 - Excellent 27-page paper on Transactional VSAM
- Key to Success
 - **Roseanna Torretto** – Session 2536
 - Goal mode conversion with over 450 performance groups
 - I was impressed with her planning and user interface
 - Key: measure before and after

This SHARE – zAAP Processors



- Great new offer from IBM
 - Can significantly reduce costs for WebSphere, DB2, IMS and CICS configurations
 - A zAAP CP is \$125K (about one-quarter the cost of a standard CP), and there are no software costs associated with it
 - You need to be very current: z890 or z990, z/OS 1.6, current releases of subsystems
 - Capacity planning and chargeback will need to be changed

This SHARE – zAAP Processors



➤ Chargeback Considerations

- E.g. – Current SMF type 30 CPU time doesn't include IFA time; instead, a new field is added for IFA CPU time
- A job could run on a zAAP engine one day and on a standard CP the next day – do you charge less for the zAAP because it's costing you less, or do you charge the customer the same in order to be consistent?
- If you charge the same, then the customers won't see any benefit from the new hardware

This SHARE – zAAP Processors



➤ Capacity Planning Considerations

- **Kathy Walsh's** recommendation: do separate reporting for each type of processor
- If you are getting a zAAP to reduce software costs, then you need to run with IFACROSSOVER=NO and force ALL Java work to the zAAP (may not provide the best performance)
- If you are getting a zAAP to simply increase current capacity at the lowest cost, then use defaults of YES for IFACROSSOVER and IFAHONORPRIORITY

This SHARE – zAAP Processors



- z890 Considerations
 - A job could run on a zAAP engine at 365 MIPS for part of a run and a standard CP at 27-365 MIPS for part of a run – how do you charge?
- Caution for reporting:
 - SMF type 30 CPU time does not include normalized IFA time, but RMF Workload Activity Report's TCB time does include normalized IFA time
- See SHARE sessions on zAAP:
 - 2892, 2867, 2561, 2825, 2514, 8366, 2878

Questions



➤ **Email:**

– cheryl@watsonwalker.com

➤ **Web site:** www.watsonwalker.com