



Cheryl's Hot Flashes #7

Cheryl Watson

Session 2543, March 8, 2002

Watson & Walker, Inc.

publishers of 'Cheryl Watson's TUNING Letter'

<http://www.watsonwalker.com>

Technology ▪ Connections ▪ Results

Cheryl's Hot Flashes

- Kudos to BMC
- User Experiences
 - FICON, SMF Synchronization, z900 High Speed Cache, RMF HTTP Report, OS/390 R10 with CICS on older CPs, LLA/VLF Situation
- APARs
 - SMF, CVAF, RLS, CA-MIM, FICON
- Interesting Items
 - ERBSCAN, XEPHON, Good Reading
- This SHARE
 - SVC Dumps, zFS Performance, WLC Defined Capacity
- WLM
 - New problem?, ERV, After Goal Mode, Release-by-release changes – which should you be considering?

Kudos to BMC

- Thanks for their comprehensive Web Site!
- Two neat examples:
 - www.dbazine.com - online DBA magazine sponsored by BMC. Excellent articles on Oracle and DB2, among others. Several by **Craig Mullins**, popular SHARE speaker.
 - Their manuals are available via PDF from their Web site. For instance, I was interested in looking at the most current CMF Monitor Batch Reference Guide for sample reports and descriptions. www.bmc.com

User Experiences

■ FICON

- **Howard Young** of American Express reports: "Just after implementation of FICON, CPU usage of IOSAS went from 3-5 seconds per one hour interval to 600 seconds per one hour interval and that usage was the same or similar 24 hours a day 7 days a week. While this did not cause any known performance problems it did increase usage of two systems in a complex by 80-100 MIPS. There is a bug in IOSVFSD in the timing of the loop control that is causing the high IOSAS utilization. The bug is expected to be fixed in APAR OW51415."
- **OW51415**: ABEND0C4 IN IOSVFSD DURING PURGEIO BECAUSE OF AN INCORRECT LOOP CONTROL VARIABLE. 02/01/14 PTF PECHANGE ORIGINATING DETAILS: ABEND0C4 in IOSVFSD during purgeio because of an incorrect loop control variable. This APAR also addresses a loop in IOSVSFD after the application of PTFs for **OW44429** and **OW44428**.

User Experiences

■ SMF Synchronization

- **Ila Miller** of University of Iowa Hospitals & Clinics reported APAR **OW47758**: SYNCVAL SYNCRONIZATION CAN BE OFF BY A SECOND DUE TO THE WAY SMF MANAGES IT'S OWN TIMER ELEMENT. I.e. SMF may write a record at 10:00:00, 10:30:00, and at 10:59:59 giving three thirty minutes intervals recorded for the 10 am period. Closed 2/28/02 as FIN, but IBM gave Ila a PTF that fixed the problem (i.e. the SMF records were written just prior to the hour).

User Experiences

■ z900 CP High Speed Cache

- **Wolfgang Schott** of IT-Austria reported the following: When ADABAS moved to z900 from a G6, the CPU time increased significantly (should have decreased). Strobe pointed out three instructions: MVI, MVC, & MVCL.
- So someone ran a test with loops of the instructions (e.g. MVI loop went from 58 seconds on a G6 to 3.01 minutes on the z900). Then they added 256 bytes between instruction and data and reran tests. CPU decreased as expected (e.g. MVI loop went to 5.27 seconds).
- Some points:
 - Loop programs where instructions and data are within 256 bytes are worthless for benchmarking, especially on the zSeries.
 - There is probably some code in ADABAS that touches data within 256 bytes of instruction that hasn't been identified or fixed yet. Will let you know.
 - Having Strobe around is mighty handy to help identify these situations!

User Experiences

■ RMF HTTP Report

- **Jerry Urbaniak** of Acxiom CDC Inc. reported problems with the RMF HTTP report. Two symptoms from RMF were: 1) Message: The skipped record contains inconsistent data. This can be caused when there are not enough records to create a report for a specific server interval. 2) Inconsistent START/END times and INTERVAL length in the HTTP report as in the following example:

```
                HTTP SERVER DETAILS
SYSTEM ID B390      START 12/11/2001-09.02.56  INTERVAL 001.30.00
RPT VERSION 02.10.00  END    12/11/2001-09.57.25
```

- IBM response: "The RMF can issue ERB482I when actual record interval is different from the time specified in SMFRecordingInterval. HTTP Server uses the log thread to write the SMF record when the SMFRecordingInterval time has reached and the log queue is full. When the Http Server is not busy, it is possible that the log queue is not full when the SMFRecordingInterval time has reached and no SMF record is written." Jerry's note: "IBM recommended the fixes for APARs **OW47050**, **OW46825**, and **PQ47542**. The first two RMF APARs were already installed. The last one is for WebSphere and is not yet applied here. Although the description for PQ47542 does not explicitly mention any changes to SMF recording, we were assured that it should correct our problems. PQ47452 will force the Webserver to cut SMF records at regular interval, and errors like ERB482I should be resolved. As we have not yet implemented the PTF we cannot comment on its effectiveness. Note that the PTF for PQ47542 is PE and also requires the fix for APAR PQ54239. There are doc changes that these PTFs introduce for WebSphere."

User Experiences

■ LLA/VLF Situation

- **Jerry Urbaniak** of Acxiom CDC Inc. reported the following: High I/O activity to the CEE.SCEERUN (or SYS1.SCEERUN), even though it was in the Linklist.
 - **OW36642** - High I/O to DASD due to LLA not staging heavily accessed modules into VLF cache causing performance problem. IBM says that the staging/destaging logic of LLA has some flaws in it that allow large modules to be destaged, even when there is high activity to them. IBM's solution is that you should move these modules into LPA. If you have high activity to your SCEERUN dataset, then you should investigate to see which modules should be forced to be resident. Argh!! (OS/390 R3+, Closed SUG 4/15/99)
 - **OW38769** - LLA may destage modules from VLF prematurely, resulting in slightly degraded performance. Opened 4/99, closed on 5/99 as FIN, included in z/OS R1 on 10/01.

User Experiences

■ LLA/VLF Situation (cont.)

- **Greg Dyck** explained why OW36642 was closed as SUG: "LLA has done its job effectively for years. The staging algorithm behind this effectiveness is extremely complex. To change the algorithm to better handle this one data set scenario carries a very real risk of introducing significant performance problems with other data sets that have, up to this point, been managed effectively. For this reason, OW36642 has been closed as SUG and no changes are planned at this time. There are other reasons besides the scenarios described by OW36642 that can cause high I/O to LLA managed libraries. The most common one is an inappropriately low value for VLF MAXVIRT, such that significant VLF trimming must occur and insufficient space exists to contain large modules. An installation can use the information contained in SMF type 41 records to perform basic tuning of the system MAXVIRT. As to the concern over what modules should be added to LPA, IBM has already identified the LE modules that it believes appropriate to be in the LPA for performance reasons, and placed these modules into SCEELPA. It is recommended that SCEELPA be made part of the LPA list. Detailed information about LE recommendations for LPA is contained in the LE Customization Manual, Chapter 7 "Placing Language Environment Modules in Link Pack, COBPACK, and LIBPACK". CEEPLPKA, the typical trouble maker for LLA, is already in SCEELPA and recommended to be placed into LPA. Other heavily hit LE modules can also be copied into SCEELPA based on the customer requirements."

User Experiences

■ LLA/VLF Situation (cont.)

- **John Zoppetti** of U.S. Steel reported: "Just after doing battle with SCEERUN, I saw the OW36642 item in your 2001 No. 5 newsletter. Our problem was that immediately following an upgrade of OS/390 R9 to R10, there were significantly higher I/O and response times to the sysres volumes, along with workload sysres delays shown in RMF Monitor III. One sysres that consistently averaged about 16 IOs/sec and 9 ms response went to about 100 IOs/sec and 29 ms response (mainly IOSQ). Also, I pinpointed I/O to the SYS1.SCEERUN DASD cylinders via the Omegamon Peek command, and Strobe measurements showed activity to CEEBINIT indicating possible repetitive LE initialization. I suspected OW36642 early on since it was similar to our problem, but did not think it applied because I added SYS1.SCEERUN to our CA PMO product and the problem continued. IBM could not tell me why the problem appeared after our OS/390 R10 upgrade but not after our R6 or R9 upgrades. They recommended either the local fix in OW36642, or using LRR or an LE Preinitialization Environment described in the LE Programmers Guide. We added SYS1.SCEELPA and modules CEEBINIT and CEEBLIBM to LPA and this resolved the problem."

User Experiences

- OS/390 R10 With CICS on Older Machines
 - **Bob Archambeault** (R. A. Solutions Int'l., Inc.) reported that some customers are seeing an increase of up to 10 times the normal SRB time when running CICS V4 and later, using transaction isolation (TRANISO=YES) migrating to OS/390 2.10 on older processors G3 or older). This was due to a change in RSM. APAR **OW52530** was closed on 2/4/02.

APARs

■ **Jerry Urbaniak** of Acxiom CDC Inc. reported the following:

- SMF Type 70 Invalid - **OW49536** - CPU partition data report uses spare/reserve CPs - D/T2604. This isn't HIPER but it's of interest to performance people because it notes a condition where the incorrect number of logical processors is stored in the type 70 records. The field SMF70BDN is used in several calculations, and it incorrectly includes the number of spare and reserve CPs. This APAR also documents some other SMF70 fields that were incorrectly documented in the OS/390 R10 manual. (OS/390 R10+, 7/6/01)
- CVAF Improvement - **OW50243** - Allow CVAF VTOC read IO (EXCP) to exploit PAVs. IBM claims a six times improvement for CVAF with this APAR. DFSMS/MVS 1.4+, 8/3/01)
- RLS Data Loss - **OW51662/OW52081/OW52080** - Records skipped reading ESDS sequentially in RLS, HIPER DATALOSS. (Note: this isn't especially performance related, however I consider it poor performance if you bypass some of your data when reading it, even though processing may take less time!) (all VSAM RLS, 12/9/01)
- There is an CA-MIM APAR **L195263** that is a list of all IBM APARs that affect the performance or operation of CA-MIM. Excellent for CA-MIM users.

APARs

■ FICON and I/O Priority

- **Brian Currah** of BDC Computer Services Inc. pointed out the following APARs:
 - **OW51126** ("WLM support to allow I/O priority management to work correctly with an ESS connected via FICON channels") describes a situation where the channel subsystem is unable to determine the actual connect time for FICON-channel-attached ESS devices. With FICON ESS devices, part of the connect time can actually be a delay due to multiplexing. The use of connect time for these devices results in an artificially low performance index (PI), thus not giving priority to jobs that need help. This APAR simply forces a 1 ms connect time for all FICON I/O requests. This time is only seen by WLM for calculating the performance index. Other measurements, such as SMF and RMF data, will still see the actual connect time. I'm not sure that I like this solution, but I guess it's the best they could do at the time. (OS/390 R8+, 10/08/01)
 - **OW47845** ("New Function APAR") is the supporting APAR for the IOS support. (OS/390 R8+, 10/08/01)

Interesting Items

■ ERBSCAN

- Introduced with RMF OS/390 R4. Allows you to browse SMF records via TSO, with some formatting for RMF records. Few people know about it, but it's VERY handy!
- Use ISPF 3.4 to display SMF data sets. Type in 'ERBSCAN' to left of data set to get list of records by timestamp and record type.
- Select an SMF record with 'ERBSHOW recnum' and it will be displayed in hex and character within record sections.
- See the RMF manuals for instructions.
- Highlighted in RMF Newsletter #13 on the RMF Web site:
www.ibm.com/servers/eserver/zseries/rmf

Interesting Items

■ Xephon New Publication

- Nifty new (and free) weekly online newsletter about MVS, DB2, MQ, CICS, etc.
- www.mainframeweek.com

■ Good Reading

- **Red Paper** – REDP0162 - DB2 for z/OS and OS/390 Version 7 Selected Performance Topics (2/19/02)

This SHARE

- Important Items From This SHARE
 - The following items were presented or brought up at this SHARE. You might have already seen them, but I thought these were important enough, or perhaps only seen by a few people, that I should bring them up again.
- Help for SVC Dumps!
 - **Jürgen Holtz** (IBM Germany) in session 2515 (What's New in z/OS WLM) made the following point:
 - z/OS V1R3 allows ASM to utilize dynamic PAV aliases if the page data set is on an ESS. (Prior to this, paging volumes weren't eligible for PAVs). A change in ASM lets it to make multiple requests to the same paging device via PAVs, so that an SVC dump shouldn't (theoretically) stop page-in requests for other work. Sadly, no benchmarks are available for this, but it's exciting! Thanks to **Greg Dyck** for implementing this. I've VERY excited about the potential of this change!

zFS Performance

- zFS Performance
 - zSeries File System – zFS
 - New feature with z/OS 1.2, but PTFs back to R10.
 - Provides significant performance improvements over HFS for USS (30% - 50% for 8k blocks). Based on DCE Local File System (Episode), but is actually a VSAM Linear Data Set. This doesn't replace HFS, but provides additional facility.
 - There are some admin items/restrictions that differ from HFS, so you don't want to jump into this, but the performance improvements should want to make you look at it.
 - Redbook – **SC24-5989** – z/OS Distributed File Service zSeries File System Administration
 - See SHARE sessions 2924 (**Jim Showalter**, IBM) & 2931 (**Peter Enrico**)


zFS Performance

- **Peter Enrico's** session 2931 - zFS & HFS Performance Considerations
 - This (and other presentations: 2562 & 2552) will be on his Web site: www.epstrategies.com
 - Peter presented some very impressive results of testing that he's done (see next foil):
 - zFS is expected to perform far superior for files 8K in size and greater and for files that are being accessed by multiple processes simultaneously
 - Peter used a relatively large file (10MB) random I/O thrasher for initial testing (since this type of file had caused him many performance problems in past with HFS)
 - Just a glimpse of we may be able to expect
 - Your results may / will vary
 - P.S. I can highly recommend Peter for consulting, on-site and public classes (WLM, WebServer, USS, performance, etc.)!

ZFS vs HFS vs TFS - Random I/Os per Sec

- Average I/Os per Second

- Thrash a 10MB File, Random I/O, 80% Reads, size of I/O 4K
- During test - no CPU or Storage contention. I/O response times within 6ms



© 2002 Enterprise Performance Strategies, Inc.



WLC Defined Capacity

- Defined Capacity

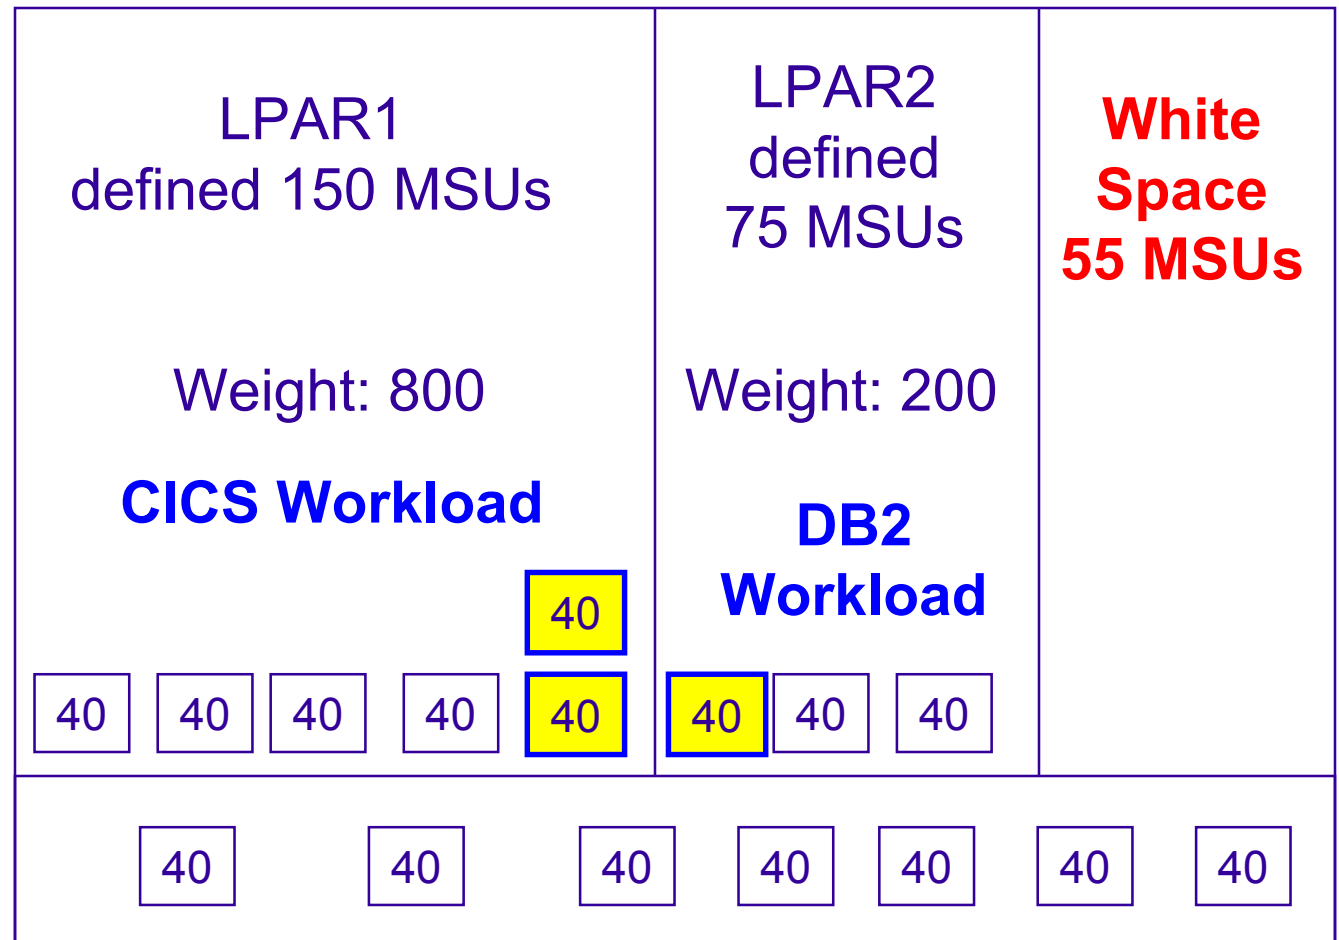
- At the last SHARE, everybody's charts showing defined capacity represented the system being soft capped to the defined capacity.
- That's not really the way it works. The system is capped based on the weight, which could be smaller or larger than the defined capacity. If weight is smaller, then the LPAR could become very constrained when the soft cap is on; if weight is larger, then weight will dynamically be adjusted down to defined cap temporarily.
- Look at presentations from **Al Sherkow** (I/S Management Strategies, Ltd.) such as session 2527, most sessions from the SAM (Software Asset Management) track (sessions 1864-1899); and **Walt Caprice's** (IBM WSC) session 2514.

WLC Defined Capacity

- **Al Sherkow's** presentations were especially interesting because of several of his examples. His presentations will be on his Web site at www.sherkow.com.
- **VWLC Price Break Point**
 - An interesting point that Al made: the pricing breakpoint for VWLC compared to PSLC is about 85%. That is, if you can set a defined capacity of 85% of your current machine, VWLC and PSLC will be about equivalent. If you define less, your VWLC will be cheaper; if you define more, your VWLC will be higher.
- **Example of Defined Capacity**
 - As one of Al's example, the next foil from his presentation shows LPAR2 with a defined capacity of 75 MSUs. But when soft cap is applied, the system doesn't use 75 MSUs, it uses the weight to calculate the cap of 56 MSUs (20% * 280 MSUs). What happens to DB2 when only 56 MSUs is available for that LPAR?
- **P.S.** I can highly recommend Al for consulting and on-site classes (capacity planning, performance, WLC, IRD, etc.)!

Soft Cap - Weights < Defined Cap

LPAR 2 Exceeding Defined Capacity zSeries 280 MSUs



- Capping enforces Initial Weights
- 20% of 280 is 56 MSUs
- So ...
 - Enforce 56 MSUs
 - Release
 - Enforce 56 MSUs
 - Release
- Monitor

WLM

- New problem? No APAR yet
 - One site recently had to move back to compat mode because the INITs were getting dispatched at too high a priority on a busy system and the more important onlines were being impacted. WLM doesn't allow the user to define or reset: *MASTER*, INIT, or WLM.
- ERV
 - Customers having problems with ENQ lockouts (where low-priority work, such as discretionary is holding an ENQ and doesn't get swapped back in for long enough to release the ENQ). The ERV value in IEAOPTxx is still valid and working. It's usually set to 500. I.e. if an address space is swapped out holding an ENQ, swap them back in until they've either released the ENQ or have used 500 service units. WSC suggestion (thanks **Kathy** and **Walt!**) is to go ahead and set it high (e.g. 50000). Make sure that it has time to release the ENQ before swapping back out again.

WLM

- You've gone to goal mode. What now?
 - Set up good reporting to identify when you're missing goals
 - Evaluate to see if you're exceeding too many goals
 - Make changes as needed due to hardware, software, and workload changes (see Peter Enrico's session 2552)
 - Re-evaluate use of CICS/IMS regions versus transaction goals
 - Re-evaluate use of discretionary work
 - Try to reduce the number of TSO periods to two (maybe three)
 - Look at changes in WLM to see what you can/should be using – set up a plan!

WLM: OS/390 R3

■ OS/390 R3

- I/O Priority Management (allows I/O DP to be set in order to make goals; but could cause erratic behavior). Suggested that you not use it until R8 or later with APAR applied.
- Application Environments (good to use)
 - IBM HTTP Server for OS/390 (IWEB) – used for scaleable server
 - DB2 V5 stored procedures (DB2)
 - WebSphere 4.0 (CB)
 - MQ/Series workflow requests (MQ)
 - SOM client object class binding requests (SOM)
- New classification qualifiers:
 - PR - Procedure name
 - PF - Perform (nifty!)
 - PFG - Perform group
- Increased report classes from 250 to 999; decreased service classes from 250 to 100

WLM: OS/390 R4 & R5

● OS/390 R4

- Scheduling environment and scheduling environment resources defined
- New classification qualifier - PRI - Priority
- Use of SYSTEM and SYSSTC service class names within classification rules. This means you no longer need to leave the STC default blank in order to place things in SYSSTC.
- WLM Batch Initiators support added – many APARs with this, but several sites are delighted with the results!
- Ability to print service policy in GML format was added

● OS/390 R5

- New subsystem added: CB (Component Broker)
 - This will later be used for WebSphere 4.0

WLM: OS/390 R6

- Added description field to classification rules and classification groups
- APAR **OW33509** added new classification qualifier - PC (process name) for MQSeries Workflow requests
- JES3 can use scheduling environments

WLM: OS/390 R6

■ Discretionary Goal Management

– Allows discretionary work to be moved above non-discretionary work if the non-discretionary work meets the following conditions:

- Is exceeding its goal
- Uses a velocity goal of 30 or less
- -or- Uses a response time goal of over one minute

(not documented...)

- Is expected to have a PI of less than .7 after the discretionary work is moved above
 - Is not in a resource group
 - Discretionary work in a resource group is not moved
- The intention is to allow discretionary to run when other work is exceeding the goals

WLM: OS/390 R6

■ Discretionary Goal Management

– BUT!

- This can cause some problems. E.g. batch production during the day has a velocity of 20, which works fine at night when the onlines are brought down. The batch production service class gets a high velocity at night and, therefore, a very low PI. After R6, in this situation, discretionary work would run higher than prod.
- There is no documented user control for this function. This has caused some people to eliminate all discretionary work (e.g. run their lowest importance work at importance=5 with a velocity of 5).
- **Jürgen Holtz** (IBM) and **Kathy Walsh** (IBM) have a solution this problem. You can define a no-op resource group (with no min or max specified). You can assign it to one or more discretionary service classes to keep them from bumping up, or you can assign it to one (or all) goal service classes to keep them from being donors. The developers are not committing to keeping this as a user interface. It's a 'work-around'.

WLM: OS/390 R7

- APAR **OW39854** on R7 and R8 changed: Dynamic Alias Management service option from NO (the default) to YES.
- MQSeries Workflow requests, previously available only with APAR **OW33509**

WLM: OS/390 R8 & R9

● OS/390 R8

■ I/O Priority Management

- APAR **OW47667** provided support to remove I/O disconnect time from consideration in the velocity calculation. This allows you to turn on I/O Priority Management. If I/O priority management is already turned on, you will need to change the velocity goals. If not on, then you should consider turning it on, which will require changing the velocity goals. I recommend using I/O Priority Management after this APAR. This also provides support for dynamic PAVs.

■ Dynamic Alias Management for ESS

- Parallel Access Volumes (PAVs)

■ JES3 can use WLM managed initiators

● OS/390 R9

■ Multisystem Enclave Support

- Designed for DB2, but probably won't be utilized until after DB2 V7

WLM: OS/390 R10

■ CICS/IMS Region vs Transaction Goals

- Prior to OS/390 R10 with APAR **OW43812** (11/16/00) applied, the moment that you assign a single transaction to a service class, every CICS (or IMS) transaction is assigned a response time goal. With the APAR, you can assign some regions to be managed with a velocity and some with transaction goals. This is terrific new support! You can also reduce the overhead of WLM searching the CICS PBs as frequently for those regions managed by velocities (e.g. test regions and those that don't have useful transactions to classify). The sampling is every .25 seconds for regions with transaction goals and 2.5 seconds for regions with region goals. One example from **Jim Petersen** of HomeSide Lending: "We 1) set all of our Production Regions to Transaction Management 2) set all of our other CICS regions to Region Management, and 3) set the CICS Transaction Management groups to CPU Critical. Our average response time per transaction has dropped from .15 sec to .06 sec."

■ 64-bit support

WLM: OS/390 R10

- Long-term storage protection
 - For address spaces in ASCH, JES, OMVS, STC, and TSO with single period, velocity goal (or a response time goal over 20 seconds)
 - For transactions in CICS and IMS, but then the region is the one managed with storage protection
 - APAR **OW43810**
- Long-term CPU protection
 - Available for all single period service classes without discretionary goals
 - If specified on CICS/IMS transactions, then the entire region is protected
 - APAR **OW43855**

WLM: OS/390 R10

- New classification qualifiers:
 - SE - Scheduling environment name
 - SSC - Subsystem collection name
 - PX - Sysplex name
 - SY - System name (but not for JES)
 - SYG - System name group

- Restriction to prevent CICS/IMS transactions from being placed in same service class with non-CICS/IMS transactions

- Restriction to prevent a lower period in a multi-period service class from having a higher importance than on an earlier period

WLM: z/OS R1

- IRD support when running on zSeries
 - LPAR CPU Management
 - Dynamic Channel Path Management
 - Channel Subsystem Priority Queuing
- Supports Defined Capacity limit on zSeries
 - RMF added support with APAR **OW46477**
 - APAR **OW49696** provides a new facility for use by RMF Monitor III and other monitoring products. A new REQLPDAT sysevent can be called to obtain performance data related to the defined capacity value from WLM.
- APAR **OW50380** - WLM CPU using samples not adjusted correctly in basic mode. Corrects APAR **OW48268**.
- APAR **OW48268** corrects several problems with WLC logic.

WLM: z/OS R1

- Doc change: From z/OS R1 Planning: WLM regarding using dynamic PAVs:
 - 11.3.1 Workload Management Considerations
 - ...
 - IMPORTANT: If you enable dynamic alias management, you must also enable I/O priority management. So you need to specify yes for both of these options on the panel. *If I/O priority management is set to no, you will get only the efficiency part of dynamic alias management and not the goal-oriented management. This means that WLM will make alias moves that minimize overall IOS queuing, but these moves will not take service class goals into consideration. (Note that you still need to be in WLM goal mode, even if you wish to use the efficiency function only.)*

WLM - z/OS R2

- IRD Enhancement
 - Includes support for Linux management (with APAR **OW50221**)
- Couple Data Set
 - Requires a CDS created by OS/390 R4 or later
- New classification qualifier:
 - SYSH (non-z/OS partitions)
- SDSF Enclave Display
 - New ENC display screen to show enclaves
- SDSF Enclave Service Class Reset
 - SDSF SPE (**PQ50025**) adds the subsystem to the ENC display and allows an operator to reset an enclave to a different service class, to quiesce an enclave, or to resume an enclave by overtyping on the ENC display
- Crypto Enhancements
 - Crypto using and delay samples collected for CAP (Cryptographic Asynchronous Processor) and PCICC (PCI Cryptographic Coprocessors)

WLM - z/OS R3

- No support for compat mode.
 - Important recommendation: get to goal mode BEFORE going to R3! The default service policy is NOT where you want to be. IBM provides an Install Definition Utility, IWMINSTL, that can install a service policy in batch. They also provide a sample service definition in SYS1.SAMPLIB(IWMSSDEF), but it will require some modifications.
- Report Class Period Enhancement
 - If only one service class is assigned to a report class, then the report class can collect all of the period, goal, and sample data that the service class does. This allows you to have a multi-period TSO service class and still get reporting for different groups of users. **Wow!**
 - Trick: if you want to get CICS/IMS transaction distributions, then assign them to a service class and report class, but set the region to REGION managed. You'll need to define a default service class as well. There will be nothing running in the service classes, but this allows the report class to pick up the goals of the unused service classes for goal calculation. Thanks to **Jürgen Holtz** (IBM) for this suggestion.

Questions



- **Email: cheryl@watsonwalker.com**
- **Web site: <http://www.watsonwalker.com>**