

# TUNING Letter

A PRACTICAL JOURNAL OF S/390 TUNING AND MEASUREMENT ADVICE

## Inside this issue...

**Many of you have asked for an article on tuning catalogs.** Michael Friske of Fidelity Investments has written an excellent tutorial for us on catalog performance that I'm sure you'll find worthwhile. It starts on page 10.

**My SHARE Trip Report** includes my "Hot Flashes #5" session, several items about z/OS and R10, and pointers to other useful sessions. Since the sessions will be on the Web about the time you receive this newsletter, I haven't included a lot of detail on the sessions, but (hopefully) enough to whet your appetite. Please see my discussion of I/O Priority Management, a technique that I have recommended against using, but which will soon have a PTF that may correct the problems. Page 26.

**SAS performance on z900s** has been a big topic recently. SAS has provided some fixes. The whole story can be found in the Cheryl's List's 50 and 51 (page 43) and in my Hot Flashes (page 26).

**IBM has recently announced z/OS R2.** The only big surprise is that it has been delayed a month (to October 26, 2001) from the expected September delivery. It's this release, when run on a z900, that has the potential of reducing your software license charges. Page 41.

**Are you using the CICS TS 1.3 Migration Utility?** If so, be sure to read the warnings on page 37.

**Although listed at the right, the January 2001 CPU Chart has already been mailed (in January),** so you should have received it by now. If not, please let us know. We just need to tie the CPU Chart to a specific issue for our CD-ROM. The 1991-2000 CD-ROMs should have all been received by now as well. You will receive an update to the Chart dated 3/16/01. See Cheryl's List #52 on page 45 for details.

## This Issue:

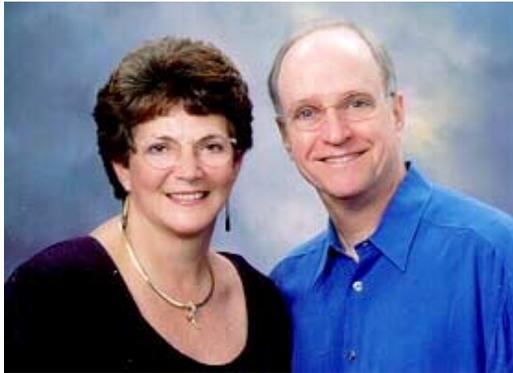
A Note from Cheryl & Tom..	2
Management Issues .....	3
S/390 News	
UIC in 64-bit .....	4
Interesting APARs .....	4
Tidbits from IBM-Main .....	5
WSC Flashes .....	5
The Net .....	5
Pubs & APARs .....	8
Focus: Tuning Catalogs	
Introduction .....	10
CAS .....	13
Catalog Sharing .....	14
Other Factors .....	16
Monitoring Catalogs .....	16
Summary .....	20
Bibliography .....	20
Notes from Cheryl .....	20
Focus: SHARE Trip Report	
Cheryl's Hot Flashes #5 .	26
R10 and z/OS .....	33
Miscellaneous .....	35
CICS Performance Tips	
Loss of DB2 Defines .....	37
What's New?	
z/OS R2 .....	41
Cheryl's List	
#49, 50, 51, 52 .....	42

## Supplement:

January 2001 CPU Chart

# A Note from Cheryl & Tom

---



We've successfully moved into our newly purchased office building at 2130 Bispham Road in Sarasota. Please make note of the address; it's in the box below. Many thanks to Linda May for performing the Herculean task of managing the entire move. It went flawlessly.

The SHARE conference in Long Beach was one of the best I've attended. The sessions were excellent, and the enthusiasm was quite infectious. z/OS seems to be well accepted. The only disturbing note was that I heard a LOT of people

say that they weren't planning to convert to WLM goal mode until they had installed OS/390 R10. R10 is the release that includes several facilities to remove the "goal mode inhibitors," as the WLM team labeled them. (I briefly described these in our 2000, No. 4 issue (page 18), but will cover them in more detail later.) I think that waiting for R10 is very short-sighted. First of all, there are hundreds of sites successfully running goal mode without the benefit of these facilities. Second, as I mention on page 31, one of the biggest jobs in goal mode migration is the conversion of report programs from performance group reporting to service class reporting. You can begin this job now, perhaps on a test or development system.

Tom and I have moved to a condo on Longboat Key, two islands north of Siesta Key. This should allow for more care-free travel, a quieter environment, and less yard work. This is our fourth home in less than two years, but I think we're finally settled!

My orchids will live in three new green houses behind our new office. Also included on the office property is a two-bedroom house. So now we can say we have a "campus," sort of like SAS! (I don't think so!)

For email, please use <[technical@watsonwalker.com](mailto:technical@watsonwalker.com)> for Cheryl, and <[marketing@watsonwalker.com](mailto:marketing@watsonwalker.com)> for Tom. ■

## Cheryl Watson's **TUNING Letter**

2001 Number 1 (Vol. 11, No. 1)

Published 6 times a year  
by Watson & Walker, Inc.  
[www.watsonwalker.com](http://www.watsonwalker.com)

Publisher: Tom Walker  
Executive Editor: Cheryl Watson  
Production Manager: Linda May

2002 SUBSCRIPTION RATES: Electronic version (CD-ROM & email) \$695 per year. Printed version \$495 per year, \$545 outside North America.

2003 SUBSCRIPTION RATES: Electronic version (CD-ROM & email) \$765 per year. Printed version \$545 per year, \$600 outside North America.

Payment may be made by a check drawn on a U.S. bank, a money order, any major credit card, or by wire transfer. Issues are mailed via first class mail. All back issues are available. Call for a complete list. Send all correspondence to: Watson & Walker, Inc., 2130 Bispham Road, Sarasota, Florida 34231, USA. Tel: 800-553-4562 or 941-924-6565. Fax: 941-924-4892. For customer service, send email to [admin@watsonwalker.com](mailto:admin@watsonwalker.com).

© 2002 Watson & Walker, Inc. All rights reserved. ISSN #1079-6606. **Reproduction of this document is permitted only for internal use at the physical address where it was received.** Permission is required for exceptions to this rule.

All IBM and other product names are trademarks of their respective owners.

NOTE: IMPLEMENTATION OF ANY SUGGESTIONS CONTAINED IN THIS NEWSLETTER SHOULD BE PRECEDED BY A CONTROLLED TEST AND IS THE RESPONSIBILITY OF THE READER.

*Please share this section with your technical management.*

## **z900 and z/OS News**

There are several items in this issue relating the new operating system, z/OS, and the new hardware, z900. The biggest news is the scare about the poor performance of some workloads on the z900, especially SAS. With quick work by IBM and SAS, a solution was found that appears to resolve the performance problems first experienced. There are a few outstanding items that are currently being worked on. See the notes about z900 status in Cheryl's Lists #50 and #51 on page 43, and my SHARE session (Cheryl's Hot Flashes #5) on page 27. My SHARE session also describes the real value of MSUs (Millions of Service Units). Since many sites are starting to use MSUs for capacity reporting, it's important to understand that the calculated value of MSU using the service unit per second value can vary from published MSUs by over 20%. There was an announcement that the hardware data compression on the z900 has significantly improved. I'm looking for confirmation of this from customers. Page 28.

z/OS R1 became orderable March 16, but several ESP sites have had a chance to use and test it. Since it's really only a slight modification of OS/390 R10, you can expect similar results. This issue lists some recent APARs (some still open) for R10 and z/OS R1. In general, it seems to be one of the most stable releases to date. See pages 4, 33, and 41. Z/OS R2 was announced with a general delivery date of October 26, 2001, a month later than expected. This was due to needed concurrent hardware updates. IBM says they expect to be able to resume their six-month schedule on March, 2002.

## **Catalog Performance**

I've received several requests for recommendations for tuning catalogs. Michael Friske of Fidelity Investments has provided an excellent article on catalog performance. Some fairly simple techniques can produce significant improvements in catalog performance. Additional APARs and notes related to catalog performance are included at the end of his article. Page 10.

## **Elsewhere in This Issue**

Our **S/390 News** on page 4 includes the following items: explanation of change in UIC from 255 to 1024 when running in 64-bit mode, interesting APARs (several still OPEN), how to keep BTAM and GAM after their withdrawal, loss of conditional JCL processing, new SMP/E function to facilitate packaging a product for shipment over a network, white papers on z/OS Firewall Technologies, a paper on IMS V7 Migration, and several new publications. In my **SHARE Trip Report** on page 26, I cover: moving to fewer CPs, DB2 archive logs, a free utility to confirm that critical duplex data sets are separate, withdrawal from service for COBOL for MVS, several WLM notes (including a new APAR that just MIGHT make I/O Priority Management worthwhile!), and pointers to some excellent SHARE sessions, which will soon be on the Web. Finally, page 37 provides a warning about using the CICS TS 1.3 Migration Facility (causing loss of DB2 Defines). ■

*Our S/390 News provides APAR updates, important announcements, recent publications and other news that we think will be of interest.*

---

## ***UIC in 64-Bit***

---

### ***New UIC Values***

When you are running either OS/390 R10 or z/OS in 64-bit mode, the UIC (Unreferenced Interval Count) value has a different range. The UIC is associated with each page frame and indicates how long the page has gone unreferenced. Thus, a value of 100 means that it's gone unreferenced for 100 seconds. The minimum, maximum, and average system UIC is kept in the SMF records. In ESA mode (31-bit), the UIC could range from 0 to 254. In ESAME mode (64-bit), the UIC can range from 0 to 2540. The mode is displayed on your RMF reports. The reason for the difference is that with so much more central storage, IBM reduced the frequency of updating the UIC to 1/10 of that previously (once per second to once per ten seconds). In order to keep the same relative meaning of the value, ten is added to the UIC count instead of one when running in 64-bit mode. Thus, a value of 100 in both modes means that the page has gone unreferenced for about 100 real (i.e. clock) seconds. If you have any programs that look at UIC, you'll need to change them before going to 64-bit mode.

---

## ***Interesting APARs***

---

### ***CPU Loop in WLM***

**Jerry Urbaniak** from Acxiom reported this Hiper APAR, [OW48081](#) - LOOP IN IRASAPRO. It documents a loop causing high CPU times in the WLM address space. The fix is applicable to R6 through R10. (3/13/01)

### ***RMF APAR***

**Jerry Urbaniak** also reported RMF APAR [OW47563](#) - Various RMF Problems (2/20/01, R3-R10). There are eleven reported problems including spin loops, abends, and invalid data in reports.

### ***R10 and z/OS APARs***

**Harold Berglas** reported the following R10 and z/OS APARs:

[OW47847](#) - ABEND0D3 IN IGG019C9 ON R2.10, **OPEN**, 1/24/01

[OW48491](#) - ABEND0C4 IN HCD MODULE CBDMUIIM AT OFFSET X'59BA' WHEN CREATING AN IODEF WITH MORE THAN 13 LPARS ON A Z900 2064-1C4 PROCESSOR, **OPEN**, 3/12/01

[OW47599](#) - MSGIEFC657I OCCURS IF A DEFINED SYMBOL IN JCL IS NOT USED D/T2064, **OPEN**, 3/9/01

- II12659** - ABEND0C4 IN OEM FDRSTART UNDER OS/390 R10 DURING SUBSYSTEM INITIALIZATION IN IEEMB860, PTF provided by vendor, 1/19/01
- OW47492** - IOS MAY NOT SET USERS REQUEST FOR USING 2K IDAWS CORRECTLY, 2/8/01
- OW47906** - IXCS1DCM FAILS VERIFYING MEMBER MST BACKED ABOVE 7FFFFFFF, 3/1/01

---

## ***Tidbits from IBM-Main***

---

IBM-Main is the most active MVS forum. Check <[bama.ua.edu/archives/ibm-main.html](http://bama.ua.edu/archives/ibm-main.html)> for archives and how to join, or use the news group at <[bit.listserv.ibm-main](mailto:bit.listserv.ibm-main)>. I've included the title of the thread so you that you can search for the entire thread in the archives.

### ***Keeping BTAM & GAM After Withdrawal From Marketing***

A user asked about the loss of BTAM after moving from OS/390 R8 to R10 (since BTAM has been withdrawn from marketing, and, therefore, from distribution). Other users referred him to APAR **II10939** (99/11/15, doc), which describes the use of the SMP/E BUILD MCS command to 'copy' the FMID from one target zone to another. But another user cautioned about possible problems (Wait state 025 during IPL!) with the technique as described in APAR **II12538** (01/02/16). With modules in two different nucleuses, IEANUC01 and IEANUC11, the technique needs to be modified, starting with R10. The new technique is pointed to by the second APAR. Thread title: os390r10.

### ***Conditional Processing***

APAR **OW46835** (DYNAMIC ALLOCATION ABEND238 OR RC1708 RSN0002 AFTER PTF FOR **OW44374**) may cause some undocumented problems. One reader found that after applying the PTF for OW46835 that conditional processing quit working. "It almost looks like it works, says 'step not run due to cond codes' but then executes the step anyway!" APAR **OW48625** (APAR OW46835) IS PE, R3-R4, 3/16/01) corrects this problem, which apparently only occurs on R3 and R4. But there is no reference to the PE on APAR OW46835. Thread title: Heads up on OW46835.

---

## ***WSC Flashes***

---

This section provides a summary of IBM's Washington Systems Center flashes. (The date of last update concludes each item.) You can get the full text at <[www.ibm.com/support/techdocs](http://www.ibm.com/support/techdocs)>.

### ***Flash10069 - Changes to Documentation Delivered with z/OS***

IBM says they're "streamlining" book shipments with this release of z/OS. I guess so. You only get three hardcopy manuals, plus one CD collection of base z/OS. All other softcopy (CD-ROM) manuals and hardcopy manuals are available for an additional fee. Everything, however, is available free from the Web. (01/25/01)

### ***Flash10070 - Supported Environments When Running z/OS as a Guest of VM***

This reiterates that fact that z/OS cannot be run as a guest of VM/ESA on a z900 machine. On a z900 machine, you can only run z/OS as 64-bit in basic mode or as a guest of z/VM running in 64-bit mode. (01/31/01)

### ***Flash10071 - SMP/E New Function - GIMZIP and GIMUNZIP***

These two new functions allow you to package a product for shipment over a network, such as a local intranet or the internet. A list of needed PTFs is supplied. Documentation can be found at <[www.ibm.com/s390/smpe](http://www.ibm.com/s390/smpe)>. The URL in the Flash is spelled wrong. (02/05/01)

### ***Flash10073 - Tivoli Storage Manager "Files with international characters stored incorrectly."***

This is a warning that TSM 4.1.2 client for Windows may not handle certain characters correctly and uncertain results may occur. APAR [IC29552](#) documents the problem. If you have downloaded or installed TSM 4.1.2, IBM recommends that you go back to 4.1.1. (02/28/01)

### ***Flash10074 - z/OS V1R1 Availability Update***

This provides notice that V1R1 is available for ordering beginning March 16, 2001. Also OS/390 R10 will remain orderable until March 2002. (03/07/01)

### ***Flash10075 - OS/390 (z/OS) SMB File/Print Automount Support***

Announcement that SMB file/print services in OS/390 R7 through R10 and z/OS R1 includes support for: automounted HFS file systems and dynamic export of HFS file systems, default permissions on a file share basis, password encryption support without requiring OCSF, and exporting MVS data sets. See APAR [OW47572](#). (03/07/01)

### ***W98007 - Problems with OPTCD=B or OPEN TYPE=J Tape Jobs and VTS***

This is an older flash describing some problems with tape jobs on VTS. This update indicates that APAR [OW33801](#) has now removed the limitation on using OPTCD=B with a VTS. (01/26/01)

### ***White Papers***

Select 'White Papers' at the WSC Web site to obtain this PDF paper:

**WP100185** - WSC WLM Migration Guide and Checklist (Note: V2 is expected by end of March).

### ***Presentations***

Select 'Presentations & Tools' at the WSC Web site to obtain these presentations:

**PRS214** - 'OS/390 Performance Hot Topics' by **Kathy Walsh** is the SHARE session by the same name. (2/28/01)

**PRS215** - 'z/OS Firewall Technologies Introduction'. (3/5/01)

**PRS216** - 'z/OS Virtual Private Network (VPN) or Better Known as Tunnels'. (3/7/01)

**PRS217** - 'Ordering Crypto for S/390 and zSeries Processors' by **Marilyn Allmond**. (3/12/01)

### ***Hits & Tips***

Select 'Hints & Tips' at the WSC Web site for the following:

**TD100408** - IMS Version 7 Migration Tips. (01/11/01)

---

## ***The Net***

---

### ***Net Locations Mentioned in This Issue***

This list includes all URLs in this issue:

Cheryl's List Archives - [www.watsonwalker.com/archives.html](http://www.watsonwalker.com/archives.html)

Amdahl SRM constants - [www.amdahl.com/doc/products/compatible/speed\\_constants](http://www.amdahl.com/doc/products/compatible/speed_constants)

IBM Announcements - [www.ibm.com/ibmlink.ibm.com/usalets](http://www.ibm.com/ibmlink.ibm.com/usalets)

IBM Catalog and VSAM Knowledge Base - [knowledge.storage.ibm.com](http://knowledge.storage.ibm.com)

IBM Catalog Performance Problems - [knowledge.storage.ibm.com/vsam/catalog-performance.shtml](http://knowledge.storage.ibm.com/vsam/catalog-performance.shtml)

IBM LSPRs - [www.s390.ibm.com/lspr](http://www.s390.ibm.com/lspr)

IBM Redbooks - [www.redbooks.ibm.com](http://www.redbooks.ibm.com)

IBM SRM constants - [www.s390.ibm.com/srm](http://www.s390.ibm.com/srm)

IBM XISOLATE - <ftp://ftp.software.ibm.com/s390/mvs/tools>

IBM WSC Flashes - [www.ibm.com/support/techdocs](http://www.ibm.com/support/techdocs)

IBM z/OS home page - [www.ibm.com/servers/eserver/zseries/zos](http://www.ibm.com/servers/eserver/zseries/zos)

IBM z/OS manuals - [www.ibm.com/servers/eserver/zseries/zos/bkserv](http://www.ibm.com/servers/eserver/zseries/zos/bkserv)

IBM-Main - [bama.ua.edu/archives/ibm-main.html](http://bama.ua.edu/archives/ibm-main.html)

SAS Service - [www.sas.com/service](http://www.sas.com/service)

SHARE - [www.share.org](http://www.share.org)

## ***Publications & APARs***

---

### ***New Redbooks***

ESA Reference Summary, SA22-7209, in PDF form:  
<[www.s390.ibm.com/ftp/books/os390/pdf/sa227209.pdf](http://www.s390.ibm.com/ftp/books/os390/pdf/sa227209.pdf)>.

### ***Technical Support Magazine***

I was surprised to find that some people were not aware of Technical Support Magazine, nor its sponsor, NaSPA (Network and Systems Professionals Association). The magazine provides very handy OS/390 (and now z/OS) technical articles. I always find articles of interest.

NASPA is a tremendous organization and provides much more than just a magazine. There are user forums, job placement services, local chapters and conferences, archives, a file library, email accounts (for a higher charge), and Internet Web hosting (also for an additional charge). The yearly charge varies from \$29.98 to \$120 depending on the services you want, but you can get the first three months free. Just go to <[www.naspa.com](http://www.naspa.com)> and sign up.

### ***IBM Hot Topics Newsletter***

The z/OS and OS/390 Hot Topics Newsletter for Spring 2001 is available for downloading. You can find it at <[www-1.ibm.com/servers/s390/os390/bkserv/hot\\_topics.html](http://www-1.ibm.com/servers/s390/os390/bkserv/hot_topics.html)>. This issue covers a lot of the new R10 and z/OS features such as IBM License Manager, Intelligent Resource Director (IRD), I/O Priority Management, and many e-business topics. This issue has a 26-page supplement on IBM documentation that's quite handy.

### ***z/OS Storage Systems***

A well written article on an important topic, "Understanding the Implications of Logical and Physical Volume Size on the Performance of z/OS Storage Systems", can be found on **Dr. Pat Artis'** Web site at <[www.perfassoc.com](http://www.perfassoc.com)>. The paper explores the performance implications of increasing logical volume and SCSI device sizes on the performance of z/OS storage subsystems. Particular emphasis is placed on the historical performance implications of larger logical volume sizes on OS/MFT and its descendents, how parallel access volumes (PAVs) address the issue for a single logical volume, potential impacts of RAID schemas on cache misses, cache sizing considerations, access density considerations, and the overall impact of SCSI device size.

### ***Items Found Elsewhere in This Issue***

#### **Manuals**

- GG24-4583 - MVS/ESA JES2 V5 Implementation
- GG66-3263 - JES2 MAS in Sysplex Environment
- SC26-4914 - Managing Catalogs
- SC33-1684 - CICS Resource Definition Guide
- SG24-2089 - OS/390 R4 Implementation
- SG24-2229 - IBM Magstar VTS: Planning, Implementing and Monitoring

SG24-4706 - MVS Parallel Sysplex Configuration Cookbook  
SG24-5151 - OS/390 R5 Implementation  
SG24-5594 - Enhanced Catalog Sharing and Management

**APARs**

This list includes all APARs and Flashes mentioned in this issue:

**Catalog** - OW39071, OW42883, OW45267, OW46456, OW38678, OW43901, OW44633,  
OW44907, OW45189, OW45267, OW45571, OW45637, OW46456, OW37739, OW45828,  
OW46165, OW47340, OW45685, II10752, OW38917, OW42831, OW43155

**DFSMS** - Flash10045, OW40167

**IOS** - OW33801

**JES2** - OW44400, OW41547, OW35410, OW46645, OW44166, OW41064, OW42769, OW43327,  
OW33073

**OS/390** - OW46835, OW44374, OW48625, OW44517, OW46338, II10939, II12538

**R10, z/OS** - OW47847, OW48491, OW47599, II12659, OW47492, OW47906, OW47572,  
OW47930

**RMF** - OW47563

**TCP/IP** - PQ41211, PQ44493, PQ45543

**TSM** - IC29552

**WLM** - OW48081, OW43091, OW47558, OW47667, OW43954

**WSC Flashes** - 10069 - 10071, 10073 - 10075, W98007 ■

# Focus: Tuning Catalogs

---

**M**ichael Friske of Fidelity Investments has been working in storage management for 15 years. He has a wide range of storage management experience including data center moves, disaster recovery planning and implementation, conversions to DFSMSHsm, SMS conversions, exploitation of new SMS functions, and more. He has also spent a lot of time working with catalogs and VSAM data sets over the past 15 years and has extensive knowledge of both.

Michael has also been actively involved with the MVS Storage Management Project at SHARE since 1987, and is currently the Project Manager. He has presented sessions on ICF catalog management and performance, VSAM management and performance, DFHSM conversions, using ABARs for disaster recovery, using DSS, PDSE experience, and most recently, eliminating JOBCAT and STEPCAT.

We've had many requests for an article on tuning catalogs. Many thanks for the article, Michael!

## Catalog Performance on Your OS/390 Systems

---

Optimizing your catalog environment for performance is not really as complex as it might appear. This article will attempt to outline the things to consider when tuning your environment. I will begin with a description of what makes up an ICF (Integrated Catalog Facility) catalog environment.

### Introduction

---

ICF catalog information is stored in two main components, the Basic Catalog Structure (BCS) and the VSAM Volume Data Set (VVDS). The BCS is usually referred to as the "catalog." It can contain information about many data sets.

The BCS is actually a specialized VSAM Key-Sequenced Data Set (KSDS). It is used to locate data sets in an OS/390 system. The BCS contains volume, ownership, and association information for cataloged data sets. It also contains the SMS (System Managed Storage) class information for SMS-managed data sets.

The VVDS is actually an extension of the Volume Table of Contents (VTOC). The VTOC is where the system maintains allocation and extent information for the data sets allocated on a disk volume. Each volume connected to an OS/390 system must have a VTOC in order to be varied online. All SMS-managed volumes and any non-SMS managed volumes with VSAM data sets on them must also contain a VVDS. A VVDS is a VSAM Entry-Sequenced Data Set (ESDS). It contains data set characteristics, extent information, SMS class information, and backward pointers to the BCS for

all types of VSAM data sets and for SMS managed non-VSAM data sets. Non-SMS, non-VSAM data sets do not have any information stored in a VVDS.

The component in the system that is used to store and access information in a BCS or a VVDS is the Catalog Address Space (CAS). CAS is a started task that is started automatically at IPL time. It processes all catalog requests on behalf of the users making the request. Each OS/390 image has a Catalog Address Space running on it.

## **Defining Catalogs for Performance**

The basic performance of a catalog environment starts with the definition of the BCS. Since the BCS is really a VSAM KSDS, all of the normal tuning considerations you would use for a VSAM data set can also be used for a BCS catalog. A catalog can be either a master catalog or a user catalog. There is one master catalog per system, which points to all of the user catalogs.

I use the following JCL to define my user catalogs:

```
//STEP010 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
  DEFINE USERCATALOG -
    (NAME (ICFCAT .APSCAT1) -
      ICFCATALOG -
      CYLINDERS (250 25) -
      FREESPACE (20 20) -
      STRNO (5) -
      BUFND (29) -
      VOLUME (CATVOL1) -
      DATA -
      (CISZ (4096))
```

A naming standard for your catalogs is very useful. This simplifies the coding of the SMS ACS routines if you place your catalogs on SMS managed volumes or if you want to exclude them from being SMS managed. Also, many DUMP/RESTORE programs have a masking capability, so naming standards help simplify your backup and recovery procedures.

Determining how much space to specify for a catalog is not an exact science. There are some general guidelines for estimating the space necessary for a catalog in the "Managing Catalogs" manual, but since catalog records are variable length records, the results of your calculating will only be a ballpark number. If you are loading a new catalog from an existing catalog, or you are re-defining an existing catalog, check the space on the existing catalog to see if it is sufficient. If it is not, make the necessary adjustments.

Many times the question is asked, how big should a catalog get before it should be split into 2 catalogs? The answer to this question, as is the case with many data processing answers, is "It depends." Having larger catalogs reduces the total number of catalogs, and therefore simplifies the management of the catalog environment (catalog placement, catalog backups, etc.). The advantages of having smaller catalogs is that fewer entries means less contention for the catalog. It also requires less time to forward recover a small catalog in the event it breaks or gets corrupted.

Most of my catalogs are about 250 cylinders big. I have a couple that are only 20 cylinders big since they have a very high numbers of accesses from some very performance sensitive applications, and I also have a couple of catalogs that are 1,000 cylinders big. These 2 catalogs contain

over 500,000 entries for data sets that are created by a report writer, and they are only accessed 2,000 – 3,000 times per day. My recommendation is to look at your catalog environment and size your catalogs so that they work in your environment.

Specifying FREESPACE is one of those things you should do if it makes sense. The existence of FREESPACE within a catalog usually reduces the number of times a Control Interval (CI) split or a Control Area (CA) split has to occur in order to insert a record. This translates into better performance. If new data sets are not added to a catalog very often, or if the inserts are not very evenly distributed, specifying FREESPACE may not be a very good idea. Very active catalogs with lots of inserts will experience a big performance benefit when FREESPACE is specified.

STRNO is a parameter that is specific to catalog definitions. This parameter is used to specify the number of concurrent read requests that a catalog can process. If it is not specified, STRNO defaults to 3. You can monitor whether this value is adequate or not using the Resource Management Facility (RMF) Monitor II or III ('ED' and 'ER' commands). To determine whether you need to increase this number, you should monitor the enqueues for the resource SYSZRPLW.catname, where SYSZRPLW is the major name of the resource and catalog name is the minor name. When I consistently see delays for this resource, and the percent delayed is greater than 30% or 40%, I consider increasing the STRNO for the catalog.

If you determine that STRNO is inadequate for a particular catalog, this parameter can be altered with the ALTER command. The values that can be specified are between 2 and 255. I find that a value of 5 works pretty well for me, even for my heavily accessed catalogs. The new value will then take affect the next time the catalog is opened. You can use the "MODIFY CATALOG,CLOSE(catname)" command from the console to close the catalog, and the next request to the catalog will cause the catalog to be opened again.

Buffer space can also be specified when a catalog is defined, or it can be altered later using the ALTER command. The default number of data buffers is STRNO+1, and the default number of index buffers is STRNO+2. I have found that the default number of index buffers usually works fine. If you have applications (such as backup jobs) that perform a lot of sequential reads, or if you have catalogs with have very long records or spanned records (like GDG records, VSAM spheres, etc.), the default number of data buffers may not be sufficient. IBM recommends that BUFND = (MAXLRECL / CISIZE) + STRNO as a minimum value. I usually specify enough data buffers for about 2 tracks worth of data.

That brings me to the CISIZE. Because most catalog requests are random, there is no big benefit to a large CI size for the data component of the catalog. A CI size of 4096 is usually a very good value. I would recommend letting the system determine the index CI size.

I let all of the rest of the parameters on a DEFINE USERCATALOG default. Using the defaults for the other parameters gives you the best performance. I especially recommend letting RECORDSIZE and SHAREOPTIONS default. The default maximum record size is 32,400. This is recommended since some types of catalog records are very large. The default SHAREOPTIONS for catalogs is (3,4). You can specify SHAREOPTIONS(3,3) only if the catalog is being placed on a volume that is defined as unshared (either using MVSCP or the Hardware Configuration Definition (HCD) facility) and will not be accessed by more than one system. If both of these conditions are not true, you should specify SHAREOPTIONS(3,4) or let them default.

## **Defining a VVDS**

If you do not define a VVDS on a volume after you initialize it, the system will define one for you the first time a VSAM data set or an SMS managed data set is allocated on the volume. The default space parameter used for this define is TRACKS(10,10). This is usually not sufficient for 3390-3 or 3390-9 volumes that are SMS managed. Once again, the "Managing Catalogs" manual contains information about estimating the space required. I usually use CYLINDERS(5 1). This has proved to be about right for most volumes in my shop. Since the VVDS can take extents, this has not been an issue.

## **The Catalog Address Space**

---

The Catalog Address Space (CAS) is started automatically at IPL time. This address space processes all catalog requests on behalf of the users making the request. These requests are normally made through a supervisor call (SVC) and passed to CAS using cross memory services.

The first time a user (TSO, online region, batch job, or started task) needs to access a catalog, CAS allocates and opens the catalog on behalf of the user. CAS then processes the request to locate, catalog, uncatalog, or alter a data set. CAS will keep the catalog opened until the next system IPL or until the "MODIFY CATALOG,CLOSE(catname)" command is issued.

The main thing you can do to enhance the performance of your catalog environment is to reduce the amount of I/O CAS has to do in order to service catalog requests. The simplest way to do this is to cache catalog records so that subsequent references to those records do not require I/O to disk.

There are two kinds of cache available for catalogs, In-Storage Catalog (ISC) cache and Catalog Data Space Cache (CDSC). ISC cache is an area set aside in main storage for the Catalog Address Space (CAS) that is used to cache catalog records. CDSC is a data space maintained by the Virtual Look-aside Facility (VLF).

Catalog records are cached in ISC or CDSC differently depending on whether the catalog is a master or user catalog. For master catalogs, all records accessed sequentially or by key are cached except for alias records. Alias records are kept in a separate table in main storage for CAS. For user catalogs, only records accessed by key are cached.<sup>1</sup>

### **In-Storage Catalog (ISC) Cache**

ISC cache is the default catalog cache. Each catalog that is opened by CAS is cached in ISC cache unless you specify that the catalog should be cached with CDSC, or unless you explicitly remove the catalog from ISC cache with the "MODIFY CATALOG" command.

Each user catalog using ISC cache is given a fixed amount of space for cached records. When a user catalog uses all of its allotted space in ISC, the least recently used record is removed from ISC to make room for the new record.

---

<sup>1</sup> IBM Managing Catalogs, "Caching Catalogs" section in Chapter 2, "Planning a Configuration of Catalogs"

The master catalog on a system is not limited to a specific amount of storage in ISC. All eligible records in the master catalog are cached in ISC as they are read.

### ***Catalog Data Space Cache***

Catalog Data Space Cache (CDSC) resides in a data space that you define in your COFVLFxx member in SYS1.PARMLIB. When VLF starts, it builds this data space and adds any catalogs that you have specified in your COFVLFxx parmlib member. Catalogs not in this parmlib member can be added later with the "MODIFY CATALOG,VLF(catname)" command.

Only one data space exists for CDSC. All catalogs eligible for CDSC use the same data space. Unlike ISC cache, catalogs are not limited to a specific amount of space to cache records in CDSC. When CDSC fills up, the oldest record is removed to make room for a new record.

### ***Selecting the Right Cache to Use***

IBM recommends that ISC cache should be used for catalogs that are not updated very frequently, and CDSC should be used for all other catalogs. The reason for this recommendation stems from how ISC cache and CDSC work.

When a catalog is using ISC cache on a particular system, and a sharing system updates any record in that catalog, CAS releases the entire ISC and creates a new ISC for the catalog. This occurs even if the record being updated by the sharing system is not cached in the ISC for this particular system.

The advantage of using ISC cache is that an additional call to VLF is not needed to obtain records that are cached. ISC cache resides in the address space for CAS. The overhead associated with a call to VLF is not really that great, but every little bit helps.

Master catalogs make good candidates for ISC cache. Usually, master catalogs are not updated very frequently, and using ISC cache will allow all records to be placed in cache as they are read. These records will not be removed from cache when other catalogs become active as they would if CDSC was being used.

Most other catalogs would benefit more from using CDSC than ISC cache. The CDSC is more intelligent than ISC caching, and it can identify individual records that a sharing system has updated. When the CDSC detects this situation, it does not usually release all of the CDSC space used for that particular catalog. Instead, it only removes the changed record from the CDSC. The entire CDSC space is only released if the number of changes by a shared system is too great to keep track of.

---

## ***Catalog Sharing Protocols and Performance***

---

It is necessary for a system making changes to a catalog to communicate those changes to all the other systems using that catalog. Traditionally, this has been accomplished using a special record in the VVDS on the volume the catalog is defined on. This record contains information about catalog records that are updated and about the current high allocated and high used RBA (record block address). Sharing systems use the information about changed records to ensure consistency of those

records if they are cached on those systems and to invalidate data and/or index buffers. The RBA information is used to update control blocks when the catalog extends beyond the current high used RBA or when it takes a new extent.

Each time a record is added, read, changed, or deleted from a catalog, I/O has to occur to the VVDS in order to read, and possibly update, the special sharing record for the catalog. A single catalog update can generate multiple I/Os to the VVDS. In environments with a high amount of catalog activity, this can impact the performance of a system or the entire sysplex. Isolating catalogs on volumes by themselves can mitigate this problem some, but this can be an expensive solution if you do not have a disk subsystem that either has virtual volumes or allows you to define volumes the size of your catalogs. In many cases, this may not even help.

Another catalog sharing protocol was introduced in DFSMS 1.5. Enhanced Catalog Sharing (ECS) mode operates very similar to VVDS mode except that the sharing information is placed in the Coupling Facility. This eliminates the I/O to the VVDS and results in better performance.

In order to use ECS mode, the following conditions must be met.

1. The SYSIGGCAS\_ECS cache structure must be defined in the Coupling Facility Resource Manager (CFRM) policy.
2. The Catalog Address Space must successfully connect to the cache structure in the Coupling Facility.
3. A catalog is referenced that meets the following criteria.
  - It must be defined with the ECSSHARING attribute set. This attribute can be altered.
  - It must be defined with SHAREOPTIONS(3,4). This attribute can also be altered.
  - It must reside on a device that is defined as SHARED.

The SYSIGGCAS\_ECS cache structure can be defined using the IXCMIAPU utility. The size of the cache structure can be specified in 256KB increments (current Coupling Facility restriction). The first 256KB increment is enough space for the fixed controls and about 40 catalogs. Each additional 256KB increment is enough space for about 60 catalogs. Specific details about sizing the SYSIGGCAS\_ECS cache structure can be found in the "Enhanced Catalog Sharing and Management" redbook from IBM.

Provided there is a physical connection from the processor to the Coupling Facility, CAS will try to connect to the SYSIGGCAS\_ECS cache structure at initialization time. Depending on how security is set up at your installation, CAS may also need to have access to the IXLSTR.SYSIGGCAS\_ECS RACF facility class.

Once CAS connects to the SYSIGGCAS\_ECS cache structure, the "MODIFY CATALOG,ECSHR(AUTOADD)" command can be used to automatically add eligible catalogs to ECS. To be eligible, catalogs must be defined with ECSSHARING and SHAREOPTIONS(3,4). These attributes can be added to an existing catalog using the IDCAMS ALTER command. All catalogs using ECS must also be defined on SHARED DASD.

Extended Catalog Sharing has been available since March 1999, but many installations have been slow to start using it. Before using ECS, I would recommend that you apply the PTF's for [OW39071](#), [OW42883](#), [OW45267](#), and [OW46456](#). These PTF's correct some major problems with ECS that could cause a major impact. Check [IBMLink](#) for details.

For more information about implementing ECS, refer to the IBM manual **Managing Catalogs (SC26-4914)**, the redbook **Enhanced Catalog Sharing and Management (SG24-5594)**, and APAR [OW38678](#).

## ***Other Factors that Affect Performance of Catalogs***

---

To serialize access to catalogs, CAS uses the SYSIGGV2 reserve. The SYSZVVDS reserve is used to serialize access to the VVDS. These two reserves are critical in order to maintain the integrity of the ICF catalog environment when multiple systems are sharing the same catalogs. To avoid a deadlock situation, which dramatically affects performance, IBM highly recommends converting the SYSIGGV2 reserve to a SYSTEMS ENQUEUE using GRS or MIM. You should not convert the SYSZVVDS reserve since it is issued with a scope of SYSTEMS. This reserve should be excluded from GRS or MIM processing.

Using JOBCAT and/or STEPCAT DD statements also has an adverse impact on catalog performance, especially if multiple catalogs are concatenated together. The system will always search the catalogs in the STEPCAT DD statement, then the catalogs in the JOBCAT DD statement, and then the catalog in the standard catalog search when locating each data set in a job step. This usually generates a lot of unnecessary I/O to the catalogs and their associated VVDS's.

In addition, catalogs specified on a STEPCAT DD statement that were not open before the beginning of the job step are closed at step termination. If these catalogs are using CDSC, this close results in the freeing of the CDSC for these catalogs. This has an immediate impact on other users on that system. Until a CDSC is rebuilt and re-populated for these catalogs, requests for these catalogs will result in physical I/Os to the volume.

IBM has been recommending the elimination of JOBCAT/STPCAT for years. To find out more about this, see my session handouts from that last Share conference entitled "How to Get Your Job Done Without Using JOBCAT/STPCAT" (session 3068).

## ***Monitoring Catalog Performance***

---

Several improvements were made in DFSMS 1.5 to give users a better idea about how their catalogs are performing. The Catalog Address Space now tracks how often it calls functions outside the catalog component and how long it takes before it gets a response. These events include calls to perform such functions as enqueues, dequeues, I/Os, SMF writes, and more. The "MODIFY CATALOG,REPORT,PERFORMANCE" command will display this information. The following is an example of the output from this command.

```

IEC359I CATALOG PERFORMANCE REPORT
*CAS*****
* -----CATALOG EVENT----- --COUNT-- ---AVERAGE--- *
* Entries to Catalog          42,888K    44.510 MSEC *
* BCS ENQ Shr Sys            66,271K    21.348 MSEC *
* BCS ENQ Excl Sys           1,214K     36.024 MSEC *
* BCS DEQ                     79,375K     0.101 MSEC *
* VVDS RESERVE CI            10,886K     0.083 MSEC *
* VVDS DEQ CI                10,886K     0.152 MSEC *
* VVDS RESERVE Shr           61,070K     0.232 MSEC *
* VVDS RESERVE Excl          834,766     0.202 MSEC *
* VVDS DEQ                    61,886K     0.194 MSEC *
* SPHERE ENQ Excl Sys        1,287K     21.378 MSEC *
* SPHERE DEQ                  1,287K     0.056 MSEC *
* CAXWA ENQ Shr                2           0.009 MSEC *
* CAXWA DEQ                     12          0.014 MSEC *
* VDSPM ENQ                   70,470K     1.288 MSEC *
* VDSPM DEQ                   70,470K     0.103 MSEC *
* RPL ENQ                      68,647      5.336 MSEC *
* RPL DEQ                      68,647      0.213 MSEC *
* BCS Get                      63,687K     0.639 MSEC *
* BCS Put                       1,012K     2.445 MSEC *
* BCS Erase                     490,669     1.992 MSEC *
* VVDS I/O                      74,656K     1.844 MSEC *
* VLF Delete Major             2,783      0.259 MSEC *
* VLF Create Minor             2,614K     0.158 MSEC *
* VLF Retrieve Minor           41,825K     0.020 MSEC *
* VLF Delete Minor             2,176K     0.103 MSEC *
* VLF Define Major              1           0.580 MSEC *
* VLF Identify                 22,053      0.019 MSEC *
* RMM Exit                     329,912     0.001 MSEC *
* Tape Exit                     329,912     0.000 MSEC *
* BCS Allocate                  8           0.113 MSEC *
* SMF Write                     4,271K     1.290 MSEC *
* VVDS Format                     37        188.267 MSEC *
* IXLCONN                       2          33.249 MSEC *
* IXLDISC                        1           0.002 MSEC *
* IXLCACHE Read                41,027      0.330 MSEC *
* IXLCACHE Write                2,746      0.177 MSEC *
* IXLVECTR                       988         0.007 MSEC *
* Resolve Symbolic              2,803      0.285 MSEC *
* MVS Allocate                   49        88.840 MSEC *
* ENQ SYSZCATS LCB              2           6.016 MSEC *
* DEQ SYSZCATS LCB              2           0.032 MSEC *
* Lookup/Pin UCB                2,582      0.025 MSEC *
* Unpin UCB                     2,582      0.019 MSEC *
* Capture UCB                   18,075K     0.017 MSEC *
* Uncapture UCB                  5,458K     0.012 MSEC *
* SMS Active Config             10,629      0.370 MSEC *
* SMS Mgmt Class ACS            1,841     486.438 MSEC *
* RACROUTE Auth                 4,730K     0.207 MSEC *
* RACROUTE Define               155,303     0.198 MSEC *
* DADSM Scratch                 362,599     52.729 MSEC *
* DADSM Rename                  2,627      90.969 MSEC *
* DADSM Allocate                221        12.413 MSEC *
* DADSM Allocate SMS            155,456     20.534 MSEC *
*CAS*****

```

The statistics that are displayed are cumulative numbers since CAS was last started or since the "MODIFY CATALOG,REPORT,PERFORMANCE(RESET)" command was issued. This information can be used to diagnose catalog performance problems. For example, if the average time for ENQs is high, perhaps you need to take a look at your GRS or MIM configuration to see if it needs to be optimized. High I/O times may be an indication that you have some channel, controller, or

volume bottlenecks. High numbers of VVDS RESERVEs, DEQs, and I/Os is an indicator that Enhanced Catalog Sharing using the Coupling Facility may benefit your shop.

Another useful command that can be used to monitor catalog performance is the "MODIFY CATALOG,REPORT,CACHE" command. The following is an example of the output from this command.

```
IEC359I CATALOG CACHE REPORT
*CAS*****
* HIT% -RECORDS- -SEARCHES --FOUND-- -DELETES- -SHR UPD- --PURGE-- *
*
* ICFCAT.APSCAT4 (VLF) *
* 98% 233 16,689 16,456 0 0 0 *
* ICFCAT.DB2PRD1 (VLF) *
* 83% 275 1,624 1,349 2 0 14 *
* ICFCAT.APSCAT5 (VLF) *
* 95% 6,106 243,404 233,576 11,542 336 42 *
* ICFCAT.SYSCAT3 (VLF) *
* 90% 1 10 9 0 0 10 *
* ICFCAT.TSTCAT1 (VLF) *
* 83% 4 37 31 0 0 345 *
* ICFCAT.APSCAT6 (VLF) *
* 91% 26 948 864 55 50 396 *
* ICFCAT.TSTCAT2 (VLF) *
* 72% 152 5,657 4,114 707 257 41 *
* ICFCAT.APSCAT1 (VLF) *
* 88% 3,760 827,100 733,014 37,748 3,481 4 *
* ICFCAT.SYSCAT5 (ISC) *
* 90% 2 22 20 0 0 4 *
* ICFCAT.SYSCAT4 (VLF) *
* 91% 1,000 16,666 15,274 395 73 182 *
* ICFCAT.TSOCAT1 (VLF) *
* 73% 4 4,145 3,060 474 265 17 *
*CAS*****
IEC352I CATALOG ADDRESS SPACE MODIFY COMMAND COMPLETED
```

This report indicates whether a catalog is using ISC cache or CDSC (indicated by "VLF" after the catalog name). It also shows the percentage of records that were found in cache when a request was made. The column this is in is labeled "HIT%". The lower this number is, the more physical I/O CAS has to do in order to get the catalog records. The IBM "Managing Catalogs" manual states that this number should be greater than 20% to make the I/O saved worth the processor storage and cycles used for the cache.<sup>2</sup>

The number of records that are currently in cache for each catalog is shown in the second column (labeled "Records") of this report. This number may not be very accurate for catalogs using CDSC due to VLF trimming, but it provides a pretty good indication about how much space is being used by a particular catalog. The next column indicates how many searches have been performed looking for records, and the fourth column shows how many records were found in the cache when these searches were performed. The number of records found divided by the number of searches gives you the HIT percent in the first column of the report.

The column labeled "DELETES" reflects the number of records that were removed from cache because of an update or a delete on this system or on any sharing system. The "SHR UPD" column reflects only the number of records that were removed from cache because of an update or a delete

<sup>2</sup> IBM Managing Catalogs, "Evaluating Catalog Data Space Cache Performance" section in Chapter 8

on another system sharing the catalog. This number is not incremented unless a record updated or deleted on another system also exists in cache on the system the "REPORT,PERFORMANCE" command was issued on.

The numbers reported in the RECORDS, SEARCHES, FOUND, DELETES, and SHR UPD columns will not always have a one-to-one correlation to the catalog functions requested by your users. This is because sometimes a single catalog request can result in the need to access multiple catalog records. You should also be aware that these numbers are reset after each IPL, after each restart of the Catalog Address Space, and a STOP/START of VLF. Because of this, the HIT percentage for each catalog will be low after one of these events until cache gets re-populated with records from these catalogs. A record does not get loaded into ISC or CDSC cache until it is accessed for the first time. It will then stay in cache until it is removed because it has not been referenced in some time and the space is needed for other catalog records or because the record was invalidated due to an update or a delete from any system sharing the catalog.

The last column in this report is labeled "PURGE". When the value in this column is greater than zero for catalogs using ISC, it means that entire ISC cache was purged for a catalog when a record in that catalog was updated on another system. If this number gets very large for a particular catalog, it is probably not a very good candidate for ISC. Each time the ISC cache is purged for a catalog, subsequent requests for records in that catalog will have to do I/O to the volume the catalog resides on until the ISC cache gets repopulated.

If a catalog is using CDSC cache, and the number in the "PURGE" column is greater than zero, it means that more updates or deletes occurred in a given time on a shared system than could be processed by the system being reviewed. Each time this occurs, the entire CDSC cache is purged for this particular catalog. Under normal conditions, only the updated record is removed from the CDSC cache, so the number in the "PURGE" column should stay small. If this number increases, it may mean that there is a problem that is preventing a system from processing updates from other systems. If this is not the case, it could mean that this catalog is not a very good candidate for CDSC.

Monitoring these statistics can provide some insight into how well cache is helping, or hurting your catalogs. It will also indicate how catalog sharing is affecting the performance of catalogs on each system. Of course, these statistics are not the whole story, but it does give you some indication about what is going on.

## ***Recommended Maintenance***

---

The following is a list of Catalog APARS I would recommend you put on your systems. Nearly all of these APARS are marked "Hiper".

**OW43901** - LEVEL SET APAR FOR HDZ11D0 CATALOG COMPONENT (DFSMS 1.4 users only)

**OW44633**- MSGIDC3009I RC50 RSN96 050-096 DURING DEFINE OF DATA SET

**OW44907** - NON-VIO AUXILIARY STORAGE IS GROWING IN SP229 FOR CATALOG ADDRESS SPACE

- OW45189** - CBSVSRT BEING USED FOR CONFLICTING PURPOSES CAUSING DATASET PROCESSING ERRORS
- OW45267** - CATALOG LOCKOUT USING ENHANCED CATALOG SHARING ON MAJOR RESOURCE SYSZCATS, MINOR = CATALOG.NAME
- OW45571** - OVERLAY IN CAS AFTER A MODIFY CATALOG COMMAND
- OW45637** - MSGIDC11705I INDEX RECORD CONTAINS DUPLICATE INDEX POINTERS AND OTHER ERRORS ON ICF CATALOGS. CATBREAKER.
- OW46456** - ECS SYSPLEX IN WAIT AFTER REBUILD EVENT GETS RC228,RSN052 IN IGG0CLSH

---

## *Summary*

---

Optimizing your catalog environment for performance is not the black art it used to be. DFSMS 1.5 provides several enhancements that allow you to monitor what is happening in your catalog environment. Using input from these new commands, you can make adjustments to your environment and then determine whether it helped. RMF also provides some help determining whether there are bottlenecks in your catalog environment.

Of course, once you optimize your catalogs, you should periodically check them to see if anything needs to be adjusted. As the workload in your systems changes, so will the access to your catalogs.

Due to space considerations, this article does not exhaust the topic of catalog performance. As such, you should reference the manuals and Web sites in the bibliography for additional details. Also, be sure to see the handouts from session 3063, "Catalog Performance Tuning and Problem Diagnosis Procedures", at the last Share Conference.

---

## *Bibliography*

---

**SC26-4914 - DFSMS Managing Catalogs**, IBM Corporation, or online at <[www.s390.ibm.com/bookmgr-cgi.bookmgr.exe/BOOKS/DGT1C111/CCONTENTS?SHELF=DGT1BK31](http://www.s390.ibm.com/bookmgr-cgi.bookmgr.exe/BOOKS/DGT1C111/CCONTENTS?SHELF=DGT1BK31)>.

**SG24-5594 - Enhanced Catalog Sharing and Management**, IBM Corporation, or online at <[www.redbooks.cim.com/redbooks/SG245594.html](http://www.redbooks.cim.com/redbooks/SG245594.html)>.

IBM Catalog and VSAM Knowledge Base web site at <[knowledge.storage.ibm.com](http://knowledge.storage.ibm.com)> or the Catalog Performance Problems web site at <[knowledge.storage.ibm.com/vsam/catalog-performance.shtml](http://knowledge.storage.ibm.com/vsam/catalog-performance.shtml)>.

---

## *Notes From Cheryl*

---

Here are some additional notes relating to catalog management:

## **Important APARs**

**Michael** mentioned the following at SHARE: See catalog/VSAM APARs for serious CAS SP252 fragmentation: [OW37739](#) (SMS 1.4, 1.5, 2/18/99), [OW45828](#) (SMS 1.3, 8/31/00), and [OW46165](#) (SMS 1.3-R10, 9/22/00). The problem, which can result in the loss of several MBs, occurs when using Jobcats, Stepcats, or when CATMAX is exceeded.

**Jerry Urbaniak** of Acxiom reported this Hiper performance problem in VLF and catalog: [OW47340](#) - LOOP HIGH CPU IN VLF MODULE COFMIDE3. CATALOG HANG, R6, R8, R10, 3/07/01.

## **CAS Performance**

**Mark Thomen** of IBM's Catalog Development team in San Jose presented session 3063 at SHARE titled "Catalog Address Space Performance and Diagnosis." It was an excellent presentation and covered some of the items mentioned in Michael's article above. Here are a couple of other notes from that session:

There have been significant changes to the catalog component starting with DFSMS 1.3. Here's an overview of what was included in each release:

- DFSMS 1.3
  - Converted code to PL/X-370 and PL/X-390
  - Rewrote using new language features
  - Eliminated obsolete and unsupported catalog features
  - Reduced pathlength and improved performance
  - Introduced new externals for sysplex and operator support
  - Improved serviceability with customer diagnosis (PD/PSI) and IBM's Service Organization (FFDC)
  - Improved performance
- DFSMS 1.4
  - Introduced a storage manager to handle most storage usage
  - Improved performance, especially for DB2 users
  - Implemented extended-direct VOLSERS to support logical SYSRES volumes
  - Improved internal diagnostic aids
  - Improved performance with shorter pathlengths
- DFSMS 1.5
  - Implemented Enhanced Catalog Sharing (ECS) to improve shared catalog performance, reduce I/O, and reduce CPU usage.
  - Extended data set aliases to simplify migration issues
  - Converted CAS code to AMODE(31)
  - Rewrote MLA table search logic
  - Raised some internal architectural limits
  - Converted most internal storage usage to above-the-line
  - Added operator commands and improved diagnostic support
  - Disabled IMBED/REPLICAT on new data sets
  - Removed password protection support

- OS/390 R10
  - Added user-managed rebuild to ECS
  - Provided automatic SVC dumps for certain categories of "user" errors (SYMREC records added via APAR [OW45685](#))
  - Recompiled on PL/X-390 to use the "relative" instruction set to improve pathlengths and performance
  - Added additional internal and external diagnostic improvements
  - Raised CATMAX to 1024 from 255 (raised default to 200 from 50)
  - Moved main load module above the 24-bit line

Mark stressed that you cannot really analyze the output of the 'Modify Catalog' commands without having a base mark to compare the output to. He suggested executing the following commands when the system is behaving properly and save the results. Then, when there seems to be a problem, execute the same commands and compare the output from the two runs. See the next page for a method of automating this function.

```
F CATALOG,REPORT
F CATALOG,LIST
F CATALOG,REPORT,CACHE
F CATALOG,REPORT,PERFORMANCE
```

The last two commands were described by Michael in his earlier article. Here is some sample output from the first two commands:

```
F CATALOG,REPORT
IEC351I CATALOG ADDRESS SPACE MODIFY COMMAND ACTIVE
IEC359I CATALOG REPORT OUTPUT 044
*CAS*****
* CATALOG COMPONENT LEVEL = HDZ11F0 *
* CATALOG ADDRESS SPACE ASN = 0022 *
* SERVICE TASK UPPER LIMIT = 180 *
* SERVICE TASK LOWER LIMIT = 60 *
* HIGHEST # SERVICE TASKS = 19 *
* CURRENT # SERVICE TASKS = 19 *
* MAXIMUM # OPEN CATALOG = 1,024 * <== This is CATMAX
* ALIAS TABLE AVAILABLE = YES *
* ALIAS LEVELS SPECIFIED = 2 *
* SYS% TO SYS1 CONVERSION = OFF *
* CAS MOTHER TASK = 009AD588 *
* CAS MODIFY TASK = 009AD3F0 *
* CAS ANALYSIS TASK = 0099DF90 *
* CAS ALLOCATION TASK = 009AD0C8 *
* VOLCAT HI-LEVEL QUALIFIER = SYS1 *
* DELETE UCAT/VVDS WARNING = ON *
*CAS*****
IEC352I CATALOG ADDRESS SPACE MODIFY COMMAND COMPLETED
```

One thing to watch out for is thrashing due to CATMAX being too low. CATMAX defines the maximum number of open catalogs that can be open at one time. If you've reached CATMAX, then one catalog must be purged and closed before you can open a new one. You don't really want to be in this position. You can use the F CATALOG,REPORT to find the current value of CATMAX, then use a 'F CATALOG,ALLOCATED', to see how many catalogs are currently allocated (you'll have to count them) and whether you need to increase CATMAX.

```

F CATALOG,LIST
IEC351I CATALOG ADDRESS SPACE MODIFY COMMAND ACTIVE
IEC347I LIST CATALOG TASK(S) 858
*****
* CAS*****
* FLAGS - TASK ADDRESS - JOBNAME / STEPNAME - ELAPSED TIME - ID *
* OW---L 0099D170 IGWLN12 / PRPX1 00.05.36 01 *
* WAITING FOR RLS Delete FROM 21C29E88 FOR 00.05.36 * hh.mm.ss
* -W-E-- 0099B190 IDAVSTAI / SFP 00.00.01 02 *
* WAITING FOR VVDS RESERVE Shr FROM 21C571BE FOR 00.00.01 *
*****
* O-OLDEST, W-WAIT, A-ABEND, E-ENQ, R-RECALL, L-RLS *
* CAS*****
IEC352I CATALOG ADDRESS SPACE MODIFY COMMAND COMPLETED

```

Mark's session also included several suggested steps for diagnosis of catalog-related problems. Information APAR [II10752](#) is useful. You can find catalog performance APARs with a keyword of CASPERFM. Mark also had an item on IBM-Main that I thought was interesting. The post was relating to a Hyper PTF from CA, but he added this comment:

*"I don't know about CA, but in our case (speaking for my components) when we mark a fix HIPER we really mean it. It is very dangerous for customers to assume that because they can't reproduce the problem from the external description that they're not going to HAVE the problem and don't need the fix.*

*"It is also my recommendation that ALL Catalog PTFs be put on - don't even BOTHER trying to sift through and decide those that may or may not be applicable to your shop."*

## Automating Catalog Data Collection

This is another gem from **Michael Friske**: "Mark Thomen talked about executing the 'Modify CATALOG,REPORT,PERFORMANCE' command before you have a problem so that you have a benchmark to compare with in the future. I have actually put together a job that uses OPS/MVS to collect this information on a daily bases and store it in a GDG data set. Here is the JCL to run the job:

```

//STEP010 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
DELETE DSM.CATSTATS.PERFORM
SET MAXCC=0
//STEP020 EXEC PGM=IKJEFT01,PARM='%FCATPERF'
//SYSPROC DD DSN=DSM.REXX.EXECS,DISP=SHR
//OPS$OPSM DD DUMMY
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD DUMMY
//STEP030 EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT1 DD DSN=DSM.CATSTATS.PERFORM,DISP=SHR
//SYSUT2 DD DSN=DSM.CATRPT.STATS.PERFORM.AM(+1),
// DISP=(NEW,CATLG,DELETE),
// DATACLAS=EXTENDED,
// SPACE=(TRK,(3,2),RLSE),
// DCB=*.SYSUT1

```

The FCATPERF REXX program referenced in STEP020 is:

```

/* REXX                                                                 */
/****-----***/
/****--- EXEC: FCATPERF          -----***/
/****--- USAGE: COLLECTS CATALOG STATS AND STORES THEM -----***/
/****---      IN A DATA SET.          -----***/
/****--- DATE: 11/01/2000          -----***/
/****--- AUTHOR: MICHAEL E. FRISKE -----***/
/****-----***/
/****---      MODIFICATION HISTORY          -----***/
/****---      -----***/
/****--- 11/01/2000 NEW EXEC          MEF -----***/
/****---      -----***/
/****-----***/

address tso
"ALLOC DA('DSM.CATSTATS.PERFORM') FILE(CATSTATS) NEW CATALOG",
"DATACLAS(EXTENDED)",
"TRACKS SPACE(3,2) DSORG(PS) RECFM(F,B) LRECL(133)"
if rc ^= 0 then
  Do
    Say 'Unable to allocate the output file!'
    EXIT rc
  End

"opscmd command('RO *ALL,F CATALOG,REPORT,PERFORMANCE')
  prefix(catresp) cmdresp(rexx) stopend(no) cmdwait(60)"
if catrespRC ^= 0 then
  Do
    Say 'OPSCMD command did not execute properly.'
    "FREE FI(CATSTATS)"
    "DELETE 'DSM.CATSTATS.PERFORM'"
    EXIT catrespRC
  End
heading = '          Statistics gathered on'
queue heading date() time()
"EXECIO 1 DISKW CATSTATS"
"EXECIO" catresp.0 "DISKW CATSTATS (FINIS STEM catresp."
if rc = 0 then
  Do
    "opscmd command('RO *ALL,F CATALOG,REPORT,PERFORMANCE(RESET)')
      prefix(catrsp) cmdresp(rexx) stopend(no) cmdwait(60)"
    if catrspRC ^= 0 then
      Do
        Say 'OPSCMD command did not properly reset the statistics'
        Say 'on all systems.'
        "FREE FI(CATSTATS)"
        EXIT catrspRC
      End
    else
      Do
        "FREE FI(CATSTATS)"
        EXIT rc
      End
    End
  End
else
  Do
    Say 'Unable to write to DDNAME=CATSTATS!'
    "FREE FI(CATSTATS)"
    EXIT RC
  End
End

```

## DFSMS 1.5 Gotchas

**Bob Shannon** and **Carl Youngren** had two catalog items in their SHARE Session 2817, Bit Bucket X'15': He used to have a "Baby" system that was a single-volume DR system. It was used to restore a full system, then it went away. It had only a master catalog, and no user catalogs. He found that ISPF option 3.4 no longer worked ("list may be incomplete") when he moved from OS/390 R5 using DFSMS 1.4 to R8 using DFSMS 1.5. The problem is described in APAR [OW38917](#) - Generic filter requests fail if MLA table is unusable. The Multi-Level Alias (MLA) table code was rewritten in DFSMS 1.5, and if there are no aliases, the MLA table is marked unusable. The solution is to allocate a user catalog and define an alias.

Aliases defined prior to DFSMS 1.5 may incorrectly use reserved bits in the catalog record. Starting with 1.5, these bits are used for SYMBOLICRELATE support. They cause SMS to attempt to relate the associated catalog name to a symbol, which causes the data sets associated with the alias to be unreachable. You get a message "IDC21364I ERROR DETECTED BY DIAGNOSE: REASON: 53 - RESERVED BITS SET IN CATALOG ALIAS ENTRY." The problem is described in APAR [OW42831](#), which is marked as HIPER. APAR [OW43155](#) adds support to Diagnose to identify which catalog alias entries have the reserved bits set. The catalog fix is to delete and redefine the alias.

**Michael Friske** of Fidelity Investments added the following suggestion to this problem. "IBM suggests deleting and defining the aliases after you put OW42831 on, but this could be a big task if you have several thousand aliases. It could also require an outage to avoid application problems during the small amount of time between the delete and the define. If you are running Catalog Solutions from Softworks, here is a job to identify and fix this problem without having to delete and redefine each alias.

"To identify the problem, run the LIST command for each catalog in your shop with the following parameters.

```
//STEP010 EXEC PGM=CSLMMAIN
//SYSIN DD *
LIST OBJECT=ICFCAT.MCAT.SYS2,
ALL,LIST=KEY,CELL=X'E7',CELL=X'03',DISP=X'5',VALUE=X'FF'
LIST OBJECT=ICFCAT.UCAT.APSCAT1,
ALL,LIST=KEY,CELL=X'E7',CELL=X'03',DISP=X'5',VALUE=X'FF'
```

"To fix the problem, run the following job with a LIST command for each alias identified in the previous job. (Remove the comments at the end of the line before running this job.)"

```
//STEP010 EXEC PGM=CSLMMAIN
//SYSIN DD *
LIST OBJECT=ICFCAT.MCAT.SYS2, <== the catalog name
KEY=C'CUST', <== the alias name
VER=(003B,FF),REP=(003B,00)
LIST OBJECT=ICFCAT.UCAT.APSCAT1,
KEY=C'CUST.PROD.LOADLIB', <== a data set alias
VER=(003B,FF),REP=(003B,00)
```

## Two SHARE Sessions

**Ron Ferguson** of Mainstar gave two excellent sessions on VSAM and catalog performance at SHARE. See sessions 3062, VSAM Tuning Basics, and 3064, Basics of ICS Catalogs, when they come out. ■

# Focus: Share Trip Report

---

*SHARE in Long Beach was excellent, as always. The sessions seemed even better than last year. I've included an expansion of my Friday morning talk, "Cheryl's Hot Flashes #5," as well as highlights from several of the sessions I attended. Most of the full presentations will be available on the SHARE Web site, <[www.share.org](http://www.share.org)>, in a month or so.*

## Cheryl's Hot Flashes #5

---

I use this session to present some of the more current information that's come across my desk.

### **MSUs Unveiled**

MSUs are the most common method of classifying a machine for software pricing. But what IS an MSU really (and how have they changed)?

'MSU' stands for 'Millions of Service Units' and is calculated as:

$$\text{MSU (rounded)} = (\text{su/sec} * \#\_of\_CPs * 3600) / 1000000$$

The 'su/sec' is defined by the hardware vendor; the #\_of\_CPs is a value from 1 to 16; '3600' is used to convert to service units per hour; and division by 1,000,000 provides the units in millions.

As an example, the 2064-102 (a 2-way) has a su/sec value of 10891.7631 and an MSU rating of 78:

$$\text{MSU (rounded)} = (10891.7631 * 2 * 3600) / 1000000 = 78$$

But lately, the calculation doesn't match the ratings issued by the vendor. (It's used for marketing and pricing, after all!) As an example, the 2064-116 has a rating of 441 MSUs, but the calculation produces 468 MSUs - a difference of over 6%. Amdahl has some even worse cases (since they calculate su/sec differently than IBM) and they show differences of up to 24%.

So if you calculate MSUs from su/sec, understand that you won't have a consistent match to the vendor's published number.

When you start to use the IBM License Manager in z/OS on a z900, it's also important to understand this difference. On the z900s, you will specify a 'defined capacity' in MSUs and the software will compare that to the actual capacity used. The software will do the calculation based on the raw su/sec of the machine. They use the su/sec for the CEC, not what the LPAR might be using. Any relationship to marketing MSUs is purely coincidental. We'll have a lot more on software pricing for z/OS on a z900 in upcoming issues.

## **z900 Status**

As I indicated in Cheryl's List #50 (see page 43), I received notes from customers indicating that their SAS runs on the z900 machines were not achieving the expected performance. Two had an under-performance of 30%; one more on the MXG forum had a 15% shortfall.

Early testing found some other programs also ran extremely poorly on the z900: Natural, the IMS log archiver DFSUARC0, and ADRDSSU (DSS). The results were up to 30% less than expected.

In the z900, the data and instruction cache are separated. The cache width is 256 bytes. If you modify something in the area that has already been pulled into the instruction cache, it must be moved back over to the data cache. The invalidation takes some overhead, and has caused increased CPU times in some programs.

The Natural program turned out to be a technique in the way a program was coded. Changing the program resolved the issue. IMS log archiver, DFSUARC0, and ADRDSSU (DSS) are currently being worked on by IBM. The theory is that some of the commonly used routines were less than 256 bytes long, so data and instructions were in the same cache.

It turned out that SAS had some prolog code that was less than 256 bytes, so that data was modified within 256 bytes of the instruction. This occurred for every function call in SAS and caused quite a bit of overhead on the z900.

SAS provided two zaps to change this logic, which (in most cases) solve the problem:

For V8: [www.sas.com/service/techsup/unotes/SN/004/004291.html](http://www.sas.com/service/techsup/unotes/SN/004/004291.html)

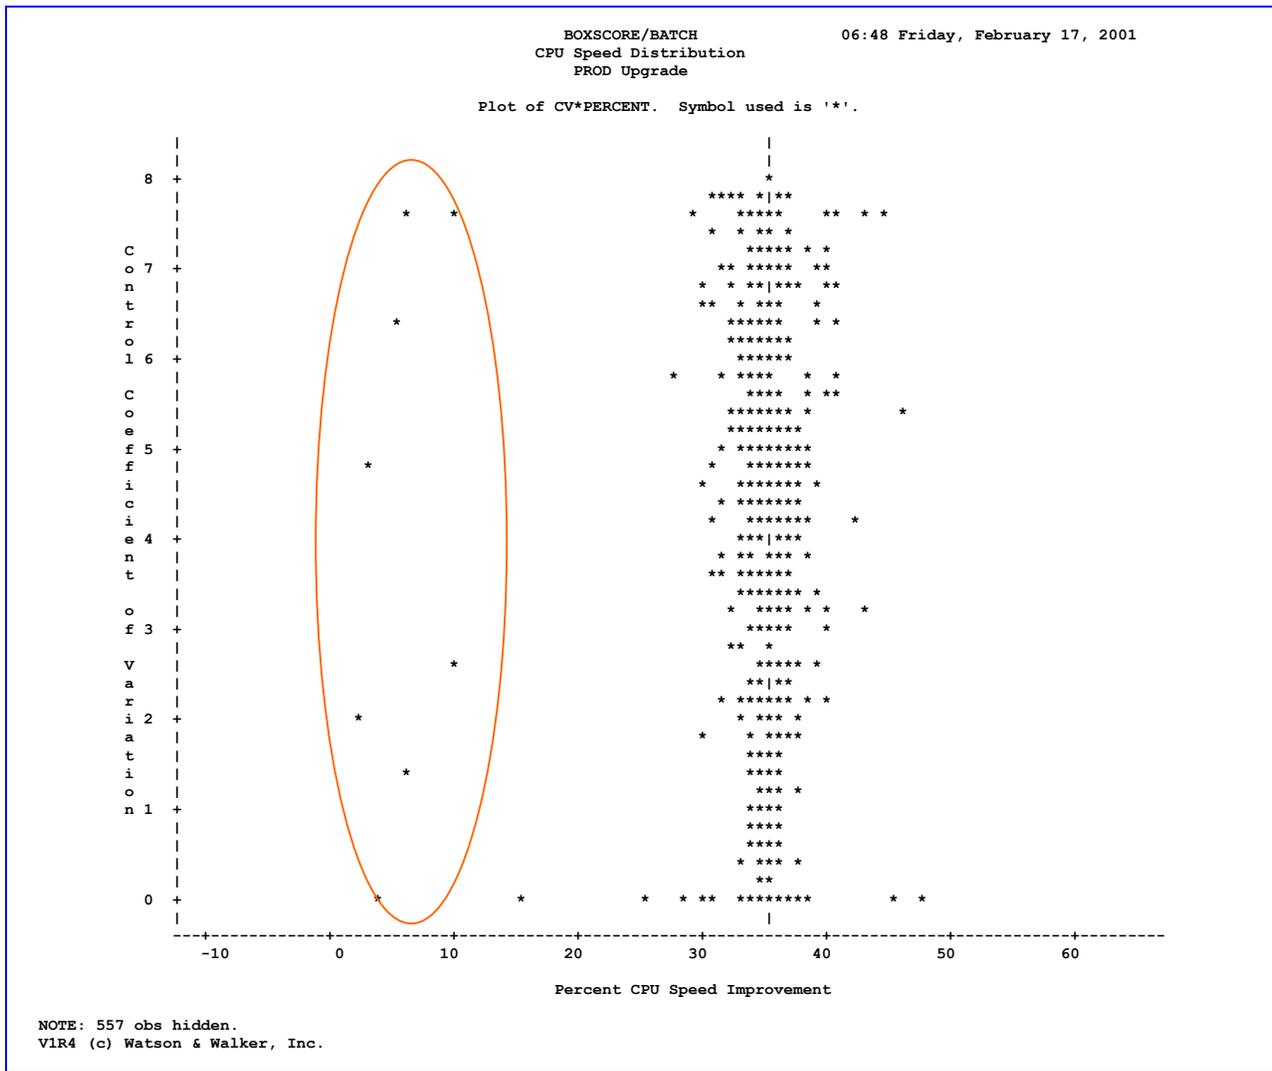
For V6: [www.sas.com/service/techsup/unotes/V6/G/G952.html](http://www.sas.com/service/techsup/unotes/V6/G/G952.html)

The good news is that these also improve SAS run times on the G5 and G6 processors by about 5%.

What all of this means to you:

- If you run SAS on a G5, G6 or z900, be sure to install the latest SAS zaps.
- If you're a vendor, consider whether your software has repetitive code that modifies itself or storage within 256 bytes of the instructions. Run tests of your product on the z900.
- When upgrading to a new processor, be sure to check each of the vendor packages on the new processor for similar situations.
- For other considerations on coding for D-cache and I-cache, see the IBM-Main archives at [<bama.ua.edu/archives/ibm-main.html>](http://bama.ua.edu/archives/ibm-main.html) and search for a string of entries with 'z/Architecture I-cache' in the heading. There were some very interesting insights during the discussion.
- Use reentrant code whenever possible. Watch out for very high use, very tiny, routines (less than 256 bytes).
- Evaluate your workloads after any processor change. See the technique described by **Al Coley** of WSC in session 2531 for determining stable and repeatable workloads (CPU per I/O and CPU per Transaction). Also consider **our BoxScore product** which identifies just this type of

**Figure 1 - BoxScore CPU Speed Distribution Report**



change after a processor upgrade. As an example, in Figure 1, each point represents a stable job step. Those points to the far left are SAS jobs running on a z900 that fell far short of their expected performance of a 37% improvement in speed. They were only getting between two and ten percent improvement.

### **z900 Compression**

Hardware compression took a major performance hit starting with G4. See our article in the 1999, No. 6 issue. But there's a change with the z900.

From IBM's z900 Ivory Letter:

*"Hardware data compression on the IBM z900 (Freeway) has been enhanced to provide significant performance gains over IBM 9672 G6 (turbo) processors. Depending upon the length of the records: a 3 to 4 times improvement has been demonstrated for compression and a 2 to 3 times improvement has been demonstrated for expansion. For both expansion and compression,*

*the longer the record length, the higher the performance gain. With expansion, records with lower compression ratios have higher performance gains".*

I'd be interested in hearing about any benchmarks on this improved compression.

### **Many CPs to Fewer, Faster, CPs**

There is a known, but long standing problem: moving work from many CPs to fewer, but faster, CPs may cause some performance and/or capacity problems. I've discussed this in several articles, the latest being in the 1998, No. 6 issue. (If you don't have this issue, it's downloadable from our Web site.) The problem occurs because some workloads, such as CICS, tend to take over the machine when moved to a faster processor. Other work on the system suffers. You could have purchased a 30% increase in capacity only to find the box maxed out on the first day.

This problem comes up once again with z900s (because of their very fast engines), but the phenomenon still exists for any similar situation.

This problem is only aggravated when people want to place too many LPARs on a single CP. Horrible things occurred to the site that moved 8 LPARs from a 6-way machine to a 2-way machine (even with 20% more MIPS). The old rule of no more LPs (logical processors) than twice the number of CPs (physical processors) still holds.

### **DB2 Archive Logs on VTS**

In my TUNING Letter (2000, No. 5) article on VTS, I passed on a comment that VTS isn't good for DB2 archive logs. Several people questioned that recommendation, since many sites use VTS for these logs. Since I couldn't remember the source for my recommendation, I polled our customers to see what they're doing. Here's the feedback from several users:

**Gary Sponseller** of University of Akron: *"We just did this a few weeks ago... It seems to work well for us. We followed the advice in Section 6.5.1 of the redbook SG24-2229-03, IBM Magstar VTS: Planning, Implementing and Monitoring."*

**Ned Diehl** of Information Systems Manager: *"If logs are large there could be recall issues in recover situations. If a physical tape must be mounted, time to first byte will be much faster with real devices than virtual ones."*

**Robert Luebke** of Sungard: *"We use it and have no major issues. We do have to account for the lack of it during D/R testing."*

**Timothy Allen** of GE Global Computer Operations: *"Logs are typically placed into a facility that is DR'able. The VTS is not DR'able."*

**Gary Stevenson** of The Boeing Company: *"We are archiving DB2 logs to virtual tape. We then have automation to copy the virtual tape to a real tape to be sent off site for disaster recovery. I think that disaster recovery would main be the reason to avoid virtual tape. I have noticed two big benefits 1) The log archive is quick - about 10 seconds for a 200MB log file. 2) When we had a tape outage, DB2 continued to archive normally; but we had to stop IMS before it ran out of log datasets. Another negative: If you run out of tapes, you can't log, which stops DB2 users in their tracks."*

**Paul Gillis** of Coles Myer: *"We are using VTS for DB2 logs without any problems, and have been doing so for about five months. We cut one log on the local VTS and the other log is cut on a remote VTS for offsite backup. SMS rules are needed to ensure the allocations via dynalloc go to the correct VTS. The only problem we have today is if the remote link is down and we cannot cut the offsite logs."*

**Collen Gordon** of Mainstar: *"There is no disaster recovery for data in a VTS except for another VTS at a remote location or backup to tape which may cause heavy thrashing in the VTS depending on when the backup is taken and if the data has been de-staged down to physical tape (and what kind of VTS you have). Many folks have found that writing 24 hours of DB2 logs to DASD, using ABARS or DFDSS to perform a backup (while on DASD) then migrating the logs off to tape using DFSMSHsm (or equivalent) provides all the coverage required and quickly gets the logs off of primary DASD when they are generally not used after being created. I would also suggest an incremental backup if you don't dual copy your offsite tapes so you can have one onsite and one offsite. Why do you need an onsite backup if you have the data set migrated? Because "stuff" sometimes gets hosed in migration. If you're running HSM audits and your environment is really clean then you probably don't need it."*

## **XISOLATE**

**Dr. Don Chesarek** from IBM developed a free tool called XISOLATE to automate the checking process to ensure that critical data sets are on different DASD.

Most key data sets have duplicate copies for availability purposes (e.g. sysplex couple data sets, JES checkpoint). These should be kept on separate hardware for best availability: different channels, different directors, different DASD, and different power. But it's often hard to keep track of where the data sets are moved. They might be allocated on different devices, but a move or migration might end up placing them on the same device. Now, there's a utility to help identify such situations.

You can use the tool to help keep data sets on different devices for availability, performance, and other business reasons.

You provide list of devices and data sets to batch job and run it at IPL and once a day. It then produces a flat file with isolation information and SYSLOG messages. XISOLATE is a non-warranted program, but code and documentation can be obtained at <ftp://ftp.software.ibm.com/s390/mvs/tools>.

Don indicates that XISOLATE has identified availability exposures in every installation that has used it during its development.

The documentation in Appendix A also provides good recommendations on which data sets you should have isolated.

## **COBOL**

As we reported in our 2000, No. 4 issue (page 27), effective March 31, 2001, IBM is discontinuing: 5668-958 VS COBOL II (Compiler, Library, Debug Facility) MVS & VM. It's replaced by -- 5688-197 COBOL for MVS & VM Release 2. (See announcement 997-070)

A less-publicized announcement is effective December 31, 2001, where IBM is withdrawing service from: 5688-197 IBM COBOL for MVS and VM. Replaced by -- 5648-A25, IBM COBOL for OS/390 and VM Version 2. (See announcement 900-220)

If you haven't started your migration, it's not too late to start!

## **More APARs**

From **Norman Hollander** of Candle Corporation: High CPU utilization can be seen in the TCPIP started task due to SNMP agent. This problem caused CPU utilization to push to 100% due to a design problem with the SNMP agent. Two fixes have been created back to 2.7. APARs **PQ41211** (R5+, 10/4/00) and **PQ44493** (R5+, 2/2/01) have fixes available. APAR **PQ45543** is closed (1/20/01) but PTFs have not been issued yet.

As we mentioned in our 2000, No. 5 issue, be sure to see WSC **Flash10045** describing APAR **OW40167** which provides a new feature for DFSMSdss. This APAR allows a new keyword, CONSOLIDATE, to consolidate multiple data extents during a DSS move.

## **WLM Notes**

Remember that compat mode disappears in 1Q02. Yea!

The EWCP sessions on WLM migration reminded me that there are still sites that need to prepare for conversion to goal mode by first getting their IPS/ICS structures organized. An article in my newsletter in 1995 addressed those issues. I've posted it on my Web site for everyone's use. See the article on positioning for goal mode at: <[www.watsonwalker.com/articles.html](http://www.watsonwalker.com/articles.html)>.

When the SHARE proceedings come out at <[www.share.org](http://www.share.org)>, be sure to look for these sessions on migration:

2510 - CICS and Workload Manager Goal Mode: New Options and Insights by **Stephen Samson**.

2511 - Migration to WLM Goal Mode by **Glenn Anderson**.

2512 - Understanding RMF Goal Mode Reports by **Cheryl Watson**.

2513 - WLM Goal Mode: Five Sane Steps on the Way to Nirvana by **Norman Hollander** has good list of APARs.

2514 - Successful WLM Migration Planning by **Jim McCoy** has good list of resources.

2515 - Taking Care of Your Loved Ones in Goal Mode - R10 Enhancements by **Jürgen Holtz**.

2516 - Managing Your Batch Workload with Throughput Manager & WLM Initiators by **Mark Jacobson**.

2539 - Advanced Goal Mode Exploitation by **Stephen Samson** describes new R10 facilities.

Users are still saying that the largest amount of time in migrating to goal mode is the time it takes to change their home-grown management reporting. I've talked to several sites who are waiting for

OS/390 R10 and its removal of goal mode inhibitors before trying goal mode. Don't wait! At least take the time to install it on your test system so you can be changing the reports at your leisure.

In R10, the new classification by system name does not apply to batch, but R10 does add classification by "jesplex" name (JES2 MAS, JES3 & whatever) via the "subsystem collection name" attribute. This solves the problem of customers with multiple jesplexes in one sysplex not being able to classify by job class (because the same job class could be defined differently in each jesplex).

Here's a compat mode APAR that you might be interested in: After adding RSU008 to an R8 system, one user had trouble meeting their SLAs. (They had 90% of TSO users completing in first period instead of 99%). APAR [OW43091](#) (Enclave CPU time is missing; R3+, 6/1/00) is in error. The new APAR is [OW47558](#) (2//2/01, OPEN).

Do you have this problem? Operators move work to different service class than intended and forget to move it back. This results in many jobs sitting in the wrong service class. How do you move them all back into the correct service class?

From **John Arwe**: You have two choices...

1. Issue a RESET RESUME on each individual job. You have to know the jobnames for this.
2. Reset all with: MODIFY WLM, MODE=COMPAT and then MODIFY WLM, MODE=GOAL. This is a little traumatic to the system and takes WLM fifteen minutes or so to get everything back to order, since everything in the system must be re-classified and monitored. If you use this technique, try to schedule it during a low period of activity (5am?).

## ***I/O Priority Management***

For me, one of the more important notices at SHARE came when Kathy Walsh of WSC announced the pending availability of a new APAR to correct the problem with I/O priority management. To help understand this, here's an introduction into I/O priority management and its problems.

When I/O priority management is turned on, the velocity calculation is changed to include I/O Using and I/O Delay and is calculated as:

$$\frac{(\text{CPU Using} + \text{I/O Using})}{(\text{CPU Delay} + \text{CPU Using} + \text{Storage Delay} + \text{I/O Using} + \text{I/O Delay})}$$

Where:

I/O Using = Connect time + disconnect time

I/O delay = Pend time + IOSQ

This made no sense to performance analysts who have always believed that disconnect time (seek time and RPS misses) is a delay.

Because of the calculation, there were several problems: velocities were quite variable; the poorer the DASD response times, the more that jobs looked like they were running better than expected and they'd get knocked down by WLM (getting the reverse effect of what's needed); the amount of disconnect time was disproportional to other delays and dominated the velocity calculation.

I always recommended that people NOT turn on I/O priority management. So did WSC. See my TUNING Letter, 2000 No. 5, pg 55.

Now with dynamic CHPID management, you MUST use I/O priority management when in goal mode for it to be effective.

So new APAR, [OW47667](#), will be available to change the calculation to omit disconnect time from the I/O USING. It's currently OPEN. Congratulations, WSC!

Apply the APAR, recalculate your velocities, and THEN you can turn on I/O Priority Management.

A slightly related APAR is [OW43954](#) - "High Disconnect Times May Occur For Devices That Present Large Device Active Times". This causes Ex Vel to show as 100%. (RMF R6-R9, 8/8/00).

See **Walt Caprice's** SHARE session 2545, WSC Performance Short Stories, for more on I/O Priority Management.

## ***R10 and z/OS***

---

### ***z/OS - The Operating System for e-business***

Session 2412 was presented by **Jerrie Stewart** and **Ed Berkel** and provided the status of OS/390 R10, z/OS R1, and z/OS R2. OS/390 V2R10 is the last release of OS/390, but will remain orderable until March 2002. It can support 64-bit real addresses when run on a z900. There is an exception to the coexistence rules to allow both R10 and z/OS R1 to coexist with R6 (n+4, an exception to the n+3 rule).

See my description of z/OS R2 on page 41.

### ***z/OS V1R1: Planning and Migrating***

**Greg Daynes** gave session 2827 for help on the migration process. Here's a brief summary of differences in base elements and optional features.

Items that are new to V1R1:

- IBM License Manager
- Msys for Setup

Items planned for removal in a future release:

- IBM Communications Server Security Level 1
- IBM Communications Server Security Level 2 LAN Server (host code only)
- LANRES
- SOMobjects ADE
- SOMobjects RTL
- Tivoli Management Framework

## **JES2**

**John Hutchinson** of WSC presented session 2652, JES2 Latest Status. Be sure to check his session handout. In the meantime, I'd like to highlight a warning and some important APARs that John mentioned.

- BERTs (Block Extension Reuses Tables) provide extensions to the JES2 Checkpoint for dynamic JQE extensions, CATs, and other control blocks. Terrible things can happen (re-IPL, cold start, etc.) if you run out of BERTs. The default (CKPTSPACE BERTNUM) is two times the JOBNUM. You can easily increase your supply of BERTs with: '\$T CKPTSPACE,BERTNUM=nnn,BERTWARN=80'. Be sure to apply the fixes for the following APARs: [OW44400](#) (BERTWARN not updated, R4-R10), [OW41547](#) (HASPHOPE not assembled). APAR [OW35410](#) is still open, so be sure to track it (Shortage w/all-member Warm start, OPEN).
- New HIPER APAR: [OW46645](#) (Spin PCE looping - Hot start required to recover, R4-R10, 1/30/01).
- APAR [OW44166](#) (Ckpt on CD data overlaid, Red Alert, R8-R10).
- Some important JES2 performance APARs: [OW41064](#) (High checkpoint writes after jobs quiesced), [OW42769](#) (Storage leaks with SDSF SWB modify), [OW43327](#) (Jobs in WLM managed class not selected).
- APAR [OW33073](#) provides some extensive JES2 IPCS enhancements to provide more checks during the dump analysis and formatting of more control blocks. As John says, "Don't wait for an emergency to install this service." This was on R4-R8, SUF 0005.

Be sure to check out his session for additional APARs and recommendations.

One of the handy items in this session was a list of deleted (but still useful) manuals, so he suggests you keep your old copies:

**SG24-2089 - OS/390 R4 Implementation**  
**SG24-5151 - OS/390 R5 Implementation**  
**GG66-3263 - JES2 MAS in Sysplex Environment**  
**SG24-4706 - MVS Parallel Sysplex Configuration Cookbook**  
**GG24-4583 - MVS/ESA JES2 V5 Implementation**

## **R10 User Experience**

**Greg Thompson** from Qwest presented session 2869, OS/390 R10 User Experience. Some interesting notes from his session:

- IEBGENER no longer works as a TSO command in R10. You must change clists to call 'sys1.linklib(iebgener)'. IBM has opened an APAR (still open), [OW47930](#), for this problem.
- They really liked the DF/HSM Fast Subsequent Migration. This feature allows a file that was previously migrated to tape and recalled, but not changed, to change pointers on the subsequent migration to simply point to the previously migrated copy. In one of their environments, they found 11,000 cases in one month, equating to a savings of 65 volumes of data, with no visible difference to the clients.

## ***New MXG Variable for ILM***

MXG change 18.298 provides a new variable, MSU4HRAV, to help you estimate your MSU usage for IBM's new License Manager. See the documentation of that change for a full description of how to use the new variable.

## ***z/OS Internet Library***

The z/OS manuals are now online at <[www.ibm.com/servers/eserver/zseries/zos/bkserv/](http://www.ibm.com/servers/eserver/zseries/zos/bkserv/)>.

---

# ***Miscellaneous Topics***

---

## ***Z/OS and OS/390 Performance "Hot" Topics***

**Kathy Walsh's** popular session 2500, WSC Performance Hot Topics, is currently available on the WSC Web site under 'Presentations'. Be sure to check out this valuable session for some very important APARs.

An important item she mentioned was the availability of a free 4.5 day seminar on IBM License Manager for qualifying customers. If you plan to run z/OS on a z900, the time spent in planning for it will be invaluable. Contact your IBM account team to see if you qualify for it.

## ***Convert IEEE Floating Point to Hex***

**Gary Hall** of IBM's WSC gave a presentation to the EWCP project on his experience in converting C++ floating point (IEEE) data to hex floating point data. As Gary found out, it isn't an easy job to locate the documentation and decipher the instructions. He ran into the need for conversion when trying to analyze the SMF type 103, subtype 2 record, used for the HTTP Server. This is the first of many fields that will be carried in IEEE floating point format. Here's the code Gary used to convert the data:

LD	R2,156(R3)	Point to C++ floating point data - LD is a floating point load instruction and R2 is a floating point register
LA	R5,564(R6)	Setup save area address
THDR	R8,R2	Convert BFP to HFP - THDR is a floating point convert instruction and R8 and R2 are floating point registers
STD	R8,0(R5)	Save HFP in the save area - STD is a floating point store instruction and R8 is the floating register used in the THDR instruction

## ***Bit Bucket X'15'***

**Bob Shannon** of John Hancock Financial Services and **Carl Youngren** of State of California Health and Human Services Data Center presented their Bit Bucket X'15, session 2817 (a great reason to stay until noon on Friday of SHARE week). Their session will be on the Web soon, but I wanted to let you know the topics ahead of time: ISPF Cut & Paste, WLM enhancements (covered in earlier newsletters), D LLA, Queued Direct I/O (QDIO) for OSA-Express Gigabit Ethernet feature, DSPCALL DSPLIST service, large tape blocksize, IEBGENER changes in OS/390 R10, re-naming a duplicate data set, fixed storage below 16 MB, HMC access via the Web, and JES2 SPOOL volume name change.

## **More Great Sessions**

These were all excellent handouts that I saw at SHARE. They'll be good to review when they appear on SHARE's Web site within a month. If they don't make it onto the Web site, I'll come back and summarize the high points in the next issue.

- 1026 - CICS TS for OS/390 Migration Planning - **Mardie Gibbs** and **Mike Jeffery**
- 1034 - CICS TS Performance - **Bob Archambeault**. Lots of detailed notes!
- 1036 - CICS and the Workload Manager - **Kathy Walsh**.
- 1040 - CICS TS and the MVS Logger - Overcoming Common Problems - **Jim Grauel**.
- 1044 - CICS TS and the MVS Logger - Performance Considerations - **Jim Grauel**.
- 1300 - What's New in the DB2 Family - **Roger Miller**.
- 1316 - 44 DB2 Tips in 60 Minutes - **Craig Mullins**. I found his home page to be an excellent source of database-related information. See <[www.craigmullins.com](http://www.craigmullins.com)>.
- 2400 - Recommended APARs - **Jerry Ng**. Full notes will be in his handout. If you want to check them out early, here they are: USS ([OW44461](#), [OW47774](#)), GRS ([OW44896](#), [OW45378](#)), XCF ([OW44648](#), [OW46877](#)), XES ([OW47217](#), [OW46531](#)), IOS ([OW45683](#), [OW47846](#)), SUP and SDUMP ([OW48081](#), [OW47171](#)).
- 2517 - Introduction to LPAR CPU Management - **Walt Caprice**.
- 2519 - Introduction to WLM Dynamic CHPID Management - **Walt Caprice**.
- 2522 - LPAR Clustering (IRD) in Action - **Joan Kelley**. This shows the results of their tests with IRD. Joan says to check the status of APAR [OW47277](#) if you're considering IRD.
- 2523 - Parallel Sysplex Tuning Update - **Joan Kelley**. While Joan gives this session every SHARE, she always adds things of importance.
- 2555 - RMF - The Latest and Greatest - **Oliver Benke**.
- 2562 - UNIX System Services and Application Tuning Issues - **Peter Enrico**.
- 2566 - S/390 Web Performance Debugging - **Judi Bank**.
- 2861 - Fully Wired Hot Topics - **Sam Knutson**. A "must-have!" It will also be available at <[www.knutson.org/share96.htm](http://www.knutson.org/share96.htm)>.
- 2866 - System Logger Procedural and Operational Considerations - **Evan Haruta** and **Nicole Burgess**.
- 2868 - OS/390 R10 User Experience at the Boeing Company - **Craig Stolze & Deborah Carnes**.
- 2877 - OS/390 in a 64-bit World - **Greg Dyck**. This session, which was previously given by Bob Rogers, presents 64-bit real considerations.
- 3000 - DFSMS R10 Technical Overview - **John Humphrey**.
- 3084 - A VTS Implementation/Migration User Experience - **Peter I. Vander Woude**. His summary: "It Rocks!". A job to process against one day's worth of SMF data used to take 8-10 minutes. With data in VTS, it now can take as little as one minute. Be careful - if the volume requested is not in the VTS cache, it can take longer than before.
- 8145 - Using SAS C/C++ Release 7.00 - **Gary Ciampa**
- 8169 - How the Linkage Editor Works: A Tutorial on Object/Load Modules, Link Editors, Loaders, and What They Do for (and to) You - **John R. Ehrman**.
- 8170 - The DFSMS/MVS Binder and Its "Program Object" Format: What The New Program Model Will Mean to You - **John R. Ehrman**.
- 8172 - eserver zSeries ("Freeway") Machines: an Assembler Programmer's View - **John R. Ehrman**.

# CICS Performance Tips

---

*René Bélanger from DaimlerChrysler Corporation sent me the following note that Pat Townsend, DB2/CICS Specialist at DaimlerChrysler sent to him. It describes an exposure for anyone using the CICS TS 1.3 Migration utility. Many thanks, René and Pat.*

## Loss of DB2 Protected & Active Thread Defines

---

Subject: Migration of Resource Control Table (RCT) to Resource Definition On-line (RDO) in CICS Transaction Server for OS/390® 1.3

When migrating the MACRO RCT to RDO in CICS Transaction Server for OS/390® 1.3, the CICS migration utility, DFHCSDUP, provides a MIGRATE command for the migration. It specifies the group name the RCT is migrated to. Entries are created in the GROUP specified on the MIGRATE command for the DB2CONN, DB2ENTRY, and DB2TRAN resource definition types. These are added to the resources defined in the CICS CSD group. The problem comes in the specific migration of parameters specified in the RCT and their conversion to CICS RDO. The definition of THRDS, THRDA, and THRDM do not migrate intact to the RDO group. There is no counterpart for THRDM at the TYPE=ENTRY level. THRDA and THRDS have counterpart definitions, but they are dropped, causing the loss of protected thread definitions and thread reuse.

The default for the equivalent parameters under RDO is 00. The parameters must be added to the DB2ENTRY define with the values for THRDA and THRDS added. The THRDA equivalent parameter is THREADLimit and the THRDS is PROTECTnum. Without the addition of these parameters on the DB2ENTRY define, our CICS system loses thread reuse with a resultant processor utilization increase for each CICS region. The resulting increase in processor utilization is dependent on the benefits of thread reuse in your environment. We experienced a processor utilization increase of 30%.

CICS system programmers may not be versed in tuning for DB2 thread reuse or mistakenly thought there was not an equivalent parameter under CICS Transaction Server for OS/390 1.3. Analysis of the RDO entries will quickly show if no values are specified for the THREADLimit and PROTECTnum parameters. The values should be set to the same values that THRDA and THRDS had been in the RCT TYPE=ENTRY define.

The following CICS Transaction Server for OS/390® 1.3 documentation was extracted from the **CICS Resource Definition Guide, SC33-1684**. Careful evaluation of the parameters of the DB2ENTRY confirms that THREADLimit and PROTECTnum provide the same RDO function as the THRDA and THRDS parameters did in the RCT. An audit of the RDO is called for.

The following manual sections should be read and understood:

Manual Section	Functional Description
1.2.1.5 CEDA and DFHCSDUP commands	This manual section describes the DFHCSDUP utility migration parameters. It is very basic and for the RCT, simply provided a group name option to hold the migrated statements.
3.2.10 MIGRATE	This manual section describes the format of the MIGRATE parameter of the DFHCSDUP utility. It provides the CSD group name that would be associated with the group list entered in the CICS SIT GRPLIST statement.  Documentation for the SIT GRPLIST parameter is found in the CICS System Definition Guide, SC33-1682, in section: 3.1.5 The system initialization parameter descriptions
5.1.4 DB2ENTRY	This manual section describes the specific parameters and their explanation for the DB2CONN, DB2ENTRY and DB2TRAN resource definition types.

~ The CICS DB2® attachment facility supplied with CICS Transaction Server is enhanced to provide RDO support for DB2 resources, as an alternative to the DSNCRCT macro. Thus the DB2CONN, DB2ENTRY, and DB2TRAN resource definition types are added to the resources defined in the CICS CSD. In this release they offer alternatives to specifying DSNCRCT macros to create a Resource Control Table (RCT), but RDO definition of these resources is expected to be the only method offered in future releases.

--- RCT macro definitions are not supported in this release (1.3), except for migration. They must be migrated to the CSD either in CICS Transaction Server for OS/390 Release 2 or in CICS Transaction Server for OS/390 Release 3. The MIGRATE command (see "MIGRATE" in topic 3.2.10) will help you to do this.

### 1.2.1.5 CEDA and DFHCSDUP commands

Table 3. CEDA and DFHCSDUP commands

Command	Function	topic
: MIGRATE	Transfers the contents of a	3.2.10
:	: terminal control table	
:	: (TCT), temporary storage	
:	: table (TST), file control	
:	: table (FCT), or destination	
:	: control table (DCT) to the	
:	: CSD. The MIGRATE command is	
:	: DFHCSDUP only.	

Section 1.2.1.5 describes the DFHCSDUP MIGRATE command. It identifies the topic section, 3.2.10, where the specific DFHCSDUP utility and MIGRATE command format is documented.

### 3.2.10 MIGRATE

The MIGRATE command will transfer the contents of a DCT, an FCT, an RCT, a TCT or a TST from a CICS load library to the CSD file. Review the following MIGRATE command format. Example II relates specifically to the RCT.

Example I. - General format of the MIGRATE command.

```
-----Migrate--Table(tablename)-----  
:                               |--Typesgroup(typesgroupname)--|  
:                               -----  
: |--TOGROUP(groupname) --|                               :  
:                               :                               :
```

#### Description

The contents of a table are transferred as one group, or as a set of several groups, containing definitions.

When migrating large tables, make sure you allocate a sufficiently large region for the largest table to be loaded.

If you are migrating a table from an earlier release of CICS, you must specify the COMPAT option on the DFHCSDUP PARM of the EXEC statement to prevent migration errors.

~ To transfer an RCT, the format is:

Example II. - Specific RCT format of the MIGRATE command.

```
MIGRATE TABLE(tablename) [TOGROUP(groupname)]
```

where TABLE(tablename) identifies the name of the table in the load library (DFHRCTxx).

The contents of a table are transferred as one group, or as a set of several groups, containing definitions. When migrating large tables, make sure you allocate a sufficiently large region for the largest table loaded. For migration purposes, RCTs must be link-edited with RMODE(24).

The result is a set of groups containing DB2CONN, DB2ENTRY and DB2TRAN resource definitions. You can define each group using the macro:

```
DSNCRCT TYPE=GROUP,GROUP=xxxxxxx
```

which you insert in the RCT source instructions before you assemble the RCT for migration. All definitions after such a TYPE=GROUP macro (up to the next TYPE=GROUP macro) go into the group named by GROUP=xxxxxxx. Definitions that occur before the first such TYPE=GROUP macro are migrated to the default group. You can also specify that definitions are to be migrated to the default group by inserting the following macro in the RCT before the definition

entries:

```
DSNCRCT TYPE=GROUP,GROUP=*DEFAULT
```

You can use the TOGROUP parameter of the MIGRATE command to assign a specific name to the default group. If you do not specify TOGROUP, the name of the default group is taken from the table name. For example, if the table name is DFHRCT24, the name of the group created is RCT24.

#### 5.1.4 DB2ENTRY

This section describes the specific DB2ENTRY parameters. Those which pertain to the definition of protected threads are shown below. The values for these should reflect values shown in the RCT before migration for THRDA and THRDS.

Review the previous RCT THRDA values to place in the THREADLimit RDO parameter.  
Review the previous RCT THRDS values to place in the PROTECTnum RDO parameter.

##### 5.1.4 DB2ENTRY

###### **PROTECTNUM**({0|value})

Specifies the maximum number of protected threads allowed for this DB2ENTRY. A thread, when it is released by a transaction and there is no other work queued, can be protected, meaning that it is not terminated immediately. A protected thread is terminated after only two complete purge cycles if it has not been reused in the meantime. Hence, if the purge cycle is set to 30 seconds, a protected thread is terminated 30 - 60 seconds after it is released, assuming it is not reused in the meantime. The first purge cycle after the CICS DB2 attachment facility has been started is 5 minutes, after which the PURGECYCLE value is applied. Threads are only protected while they are inactive. If a transaction reuses a protected thread, the thread becomes active, and the current number of protected threads is decremented.

###### **THREADLIMIT**({0|value})

Specifies the maximum number of threads for this DB2ENTRY that the CICS DB2 attachment allows active before requests are made to wait, are abended, or diverted to the pool.

###### **THREADWAIT**({POOL|YES|NO})

Specifies whether or not transactions should wait for a DB2ENTRY thread, be abended, or overflow to the pool should the number of active DB2ENTRY threads reach the **THREADLimit** number.



# What's New?

---

*The biggest news this month was the announcement of z/OS R2. Here's a brief summary, but expect more in future issues.*

## z/OS R2

---

z/OS R2 was announced on 2/27/01 in announcement letter 201-044 <[www.ibm.com/ibmlink/usalets](http://www.ibm.com/ibmlink/usalets)>. R2 will be available October 26, 2001. This is a month later than normally expected, but IBM says that they will try to get back on the 6-month schedule by March 2002. The delay is due to some expected hardware enhancements currently projected for 4Q01. z/OS R2 will run on z900s, the 9672 G5 and G6s, and the Multiprise 3000. z/OS R2 will revert back to n+3 coexistence, so can only coexist back to R8.

The main features of R2 support e-business. This includes an easier method of managing Kerberos security and more use of digital certificates. R2 will provide full support for the IBM License Manager to allow support for sub-CEC capacity licenses. This support will also become available for z/OS R1 on 9/30/01. (IBM uses that date, although I suspect that it might change to the R2 availability date of 10/26/01.) **As mentioned MANY times before, R2 is the last release to support WLM compatibility mode.**

RMF enhancements for R2 include the following: Crypto facility usage and delays, response time distributions for report classes, report support for CF structure duplexing, real-time view of each LPAR's CPU usage in MSUs for use with IBM License Manager. A minor, but frequently asked for item is support for wildcards for GRS in Resource Name lists.

The following items are announced with R2, but not expected to be available until after GA

- C++ support of ASCII and ANSI '98 standards.
- One of the highlights of R2 is support for System Managed CF Structure Duplexing. This is only available on the associated hardware, which won't be available until 4Q01. The purpose of this feature, of course, is to provide higher availability in a parallel sysplex environment.
- The Intelligent Resource Director, IRD, is enhanced to allow CPU resource direction to Linux partitions.
- Support for HiperSockets.

The z/OS home page is <[www.ibm.com/servers/eserver/zseries/zos](http://www.ibm.com/servers/eserver/zseries/zos)>.

For more information, see Session 2412 on the SHARE Web site in a month. This was a session on z/OS by **Jerrie Stewart** and **Ed Berkel**. ■

# Cheryl's List

---

*Here's a summary of the last transmission sent to subscribers of our free electronic **Cheryl's List**. We've eliminated some of the sections if they were printed in a previous newsletter. Past issues of Cheryl's List can be obtained in full at <[www.watsonwalker.com/archives.html](http://www.watsonwalker.com/archives.html)>. See the last paragraph below for instructions on how to sign up for the list.*

## Cheryl's List #49 - 31 January 2001

---

1. zSeries LSPR Update
2. CPU Chart in Spreadsheet Format!
3. Cheryl Watson's TUNING Letter 2000, No. 6 Summary (*Previously printed in TL 2000, No. 6*)
4. SHARE in Long Beach

### **1. zSeries Update**

On February 24, IBM updated the LSPR ratings previously published in December. They made the decision to update the ratings after further testing of the new processors. Models 2064-101 to 2064-109 have new performance ratings, new MSUs, and new SU/Second values. These models were downgraded from 0 to 6% from the December ratings. IBM also changed the base model from the 101 (a 12-PU MCM base) to the 1C1 (a 20-PU base). The 20-PU based models (110-116 and 1C1-1C9) were not affected from an MSU point of view. Current z900 customers can order a microcode update (Patch Stream # H25117, MCL # 040) to install the new SU/Second values.

Our January 2001 CPU Chart reflects these new values and will be sent to subscribers later this week. IBM's new LSPRs and MSUs can be found at <<http://www.s390.ibm.com/lspr>> and the SUs/Second can be found at <<http://www.s390.ibm.com/srm>>.

### **2. CPU Chart in Spreadsheet Format!**

We're pleased to announce the availability of our CPU Chart in spreadsheet format. See our description below. This is also available for non-subscribers for a fee. Go to <<http://www.watsonwalker.com/chart.html>> for more information.

### **4. SHARE in Long Beach**

My favorite user conference, SHARE, will be held in Long Beach next month. I'll be there to present two sessions: "Understanding RMF Goal Mode Reports" (#2512) and "Cheryl's Hot Flashes Part V" (#2543). Here's a description from SHARE:

SHARE, the original IBM users group, presents SHARE in Long Beach (CA) from February 25 - March 2. The event features SHARE's traditionally robust agenda of more than 1,000 hours of

technical education geared towards users of large, complex IBM systems. Highlights in Long Beach include:

- \* The latest information on IBM's new zSeries
- \* Extensive offerings on Software Asset Management
- \* Scores of user experiences
- \* In-depth Windows 2000 insights

The event also features a comprehensive products & services exhibition, hands-on labs, nightly networking events, and a complimentary Internet Café.

Complete event information - including the latest schedule of events, session abstracts, and secure online housing and registration - can be found at <<http://www.shareinlongbeach.org>>.

---

## **Cheryl's List #50 - 9 February 2001**

---

1. zSeries Update
2. January 2001 CPU Chart is Available
3. 1991-2000 CD-ROMs Mailed

### **1. zSeries Update**

I've received feedback from two z900 sites who did not receive their expected performance when running SAS-based work (up to 30% less capacity than expected). Another user reported a 13% under-capacity reading on SAS work on the MXG list forum. At least one Natural user is experiencing the same problem.

The current surmise, although not confirmed by any parties, is that the change in the hardware design of the z900s is affecting 4GL languages. The design change was to separate the Data cache and Instruction cache for each CP. This makes for faster processing of instructions and data, but can be counter-productive for any application that intersperses data and instructions or that creates instructions on the fly in a data area and attempts to execute them. The additional locking and resultant movement of instructions from the data cache to the instruction cache apparently increases the CPU time for these applications. IBM and SAS are looking into the problem in more detail at this time. Other 4GL software vendors should look into this situation as well.

At this point we don't know if all SAS programs will be affected or if these are isolated cases. But we'd love to know! If you've run some tests on your z900 (especially SAS tests), please let us know the results. Send them to <[technical@watsonwalker.com](mailto:technical@watsonwalker.com)>. Your name and company can remain confidential. (For a tool to measure the performance of a new processor or software change, don't forget our software product, BoxScore <<http://www.watsonwalker.com/boxscore.html>>.)

### **2. January 2001 CPU Chart is Available**

The PDF and Excel version of our January 2001 CPU Chart was sent to electronic TUNING Letter subscribers on February 2<sup>nd</sup>. The print version is being mailed today. This chart contains the z900

processors and MIPS based on the OS/390 R10 LSPRs. Additions to the chart include the STSI models and Architectural Level Sets (ALS).

Here are the file names:

Jan2001.xls - Excel 97 format of the chart.

Jan2001.pdf - The first 12 pages of the chart linkable from the xls file.

CW20011C.pdf - The entire 36 page chart suitable for printing.

Although Lotus users can read the xls chart, there are some incompatibilities. If you did not receive the note listing the differences, please let us know.

As an alternative, **Philip Foster** of La Confederation des Caisses Populaires et d'Economie Desjardins du Quebec (CCPEDQ) recommended downloading a free Excel 97/2000 viewer from the Microsoft Web site: <<http://officeupdate.microsoft.com/2000/downloadDetails/xlviewer.htm>>.

Speaking of Excel 2000, **Norman Hollander** of Candle pointed out that Excel 2000 doesn't translate the column comments well. In order to read more than two lines of the comments, you have to move the spreadsheet to an unprotected document.

### **3. 1991-2000 CD-ROMs Mailed**

Yesterday, the latest CD-ROMs were mailed to all electronic subscribers. This CD-ROM contains all of our issues through 2000 No. 6, the November 2000 CPU Chart in PDF and Excel format, and an additional 42 papers and presentations by me. Electronic subscribers will also receive any presentation I make within a week of its presentation; so if you can't come to SHARE or CMG, you can still see my presentations on a timely basis.

## ***Cheryl's List #51 - 9 March 2001***

---

1. zSeries Update
2. Update to High CPU APAR

### **1. zSeries Update**

In the last Cheryl's List (#50), I discussed a problem with some types of work that did not achieve their expected capacity when running on a z900 processor. Fortunately, great strides have been made since then.

The z900 design change to separate the Data cache and Instruction cache for each CP affected some programs, some quite significantly. The cache line size is 256 bytes. If you modify something within a line that has already been pulled into the instruction cache, it must be moved back over to the data cache. The invalidation takes some overhead.

The last Cheryl's List mentioned a problem with a Natural program. It turned out to be the way the program was coded. Changing the program resolved the issue. IMS log archiver, DFSUARC0, and ADRSSU (DSS) also had problems and are currently being investigated by IBM.

The most problems reported were with SAS. SAS and IBM worked closely on the problem to resolve it. It seems that SAS had some prolog code that was less than 256 bytes where data was modified within a cache line that also contained instructions. This occurred for every function call in SAS and caused quite a bit of overhead on the z900. SAS provided two zaps to change this logic, that, in most cases, completely solved the problem:

For V8: [www.sas.com/service/techsup/unotes/SN/004/004291.html](http://www.sas.com/service/techsup/unotes/SN/004/004291.html)

For V6: [www.sas.com/service/techsup/unotes/V6/G/G952.html](http://www.sas.com/service/techsup/unotes/V6/G/G952.html)

The good news is that these also improve SAS run times on the G5 and G6 processors by about 5%. It might improve SAS on the MP3000, but no tests have been run at this time.

I said it solved the problem "in most cases" because two sites still tell me that their test programs were improved after the zaps, but not by as much as expected. But, since these were both very small loop programs, I'm not sure that they're indicative of typical SAS workloads. If you see any unexpected results after moving to the z900 processors, be sure to notify IBM or SAS (and me!).

For your own programming, use reentrant code whenever possible, and watch out for very high use, very tiny, non-reentrant routines (less than 256 bytes). For an interesting discussion of considerations on coding for D-cache and I-cache, see the IBM-Main archives at <http://bama.ua.edu/archives/ibm-main.html> and search for a string of entries with 'z/Architecture I-cache' in the heading.

(As always, for a tool to measure the performance of a new processor or software change, don't forget our software product, BoxScore <http://www.watsonwalker.com/boxscore.html>. It can tell you exactly which programs are not meeting expectations, and those that are exceeding expectations as well.)

## **2. Update to High CPU APAR**

In Cheryl's List #45 (July), I discussed APAR [OW44517](#) (After OW39132, IEAVEOR (IOC003) may not decrement ASCBTCBS marked as HIPER (07/17/01)). This related to performance degradation and/or high LPAR management time (33 MIPS to 69 MIPS in uncaptured time for one site). Many sites ran into the problem and applied the fix for the APAR. **Kathy Walsh** in session 2500 at SHARE pointed out that APAR [OW46338](#) is a PE of OW44517. So if you applied OW44517, then also apply OW46338 (R5+, 1/2/01).

---

# **Cheryl's List #52 -19 March 2001**

---

## **1. Updates to CPU Chart**

There are two updates to our January 2001 CPU Chart. We will be sending out updates to the CPU Chart when we mail our TUNING Letter 2001, No. 1 later this week.

Amdahl updated the service unit per second value for their three Omniflex servers a few months ago, but their Web site wasn't updated until February. See [http://www.amdahl.com/doc/products/compatible/speed\\_constants/](http://www.amdahl.com/doc/products/compatible/speed_constants/) for their updated list.

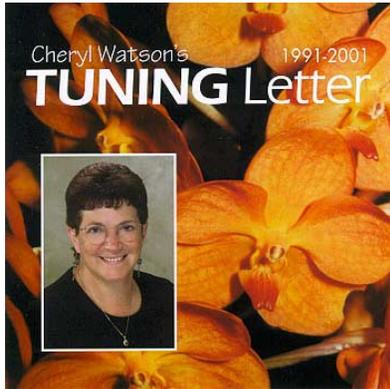
Thanks to **Svend Kjaer** of DMdata Denmark for pointing out that when I changed the CPU Chart to carry the updated MIPS for the 2064 series, I forgot to make all the necessary changes. The uni service unit per second rate for the 101 to 109 models needed to be changed and the MP ratio for models 110 to 116 needed to be based on the 1C1 instead of the 101.

---

### ***SUBSCRIBE TO CHERYL'S LIST***

If you would like to obtain an email subscription to Cheryl's List, just go to our Web page [www.watsonwalker.com](http://www.watsonwalker.com) and fill out the form under "Cheryl's List." That signs you up. Remember, it's a one-way list, from us to you. If you make a "reply", it will come just to us, not to the other members of the list. To unsubscribe, send an email message with only the word "unsubscribe" as the body of the message (no taglines!) to [cheryls-list-request@xmission.com](mailto:cheryls-list-request@xmission.com). ■

# Now via CD-ROM and EMAIL!



## Cheryl Watson's **TUNING** Letter

- ALL THE BACK ISSUES of the TUNING Letter
- A YEAR'S SUBSCRIPTION TO FUTURE ISSUES
- ALL OF IT SEARCHABLE
- ALL in ELECTRONIC FORMAT
- On CD-ROM and sent via EMAIL
- ALL FOR \$695 per year

**Order now - Prices will increase in 2003**

**Subscribe now** and we'll send you a **CD-ROM** containing **every back issue** of Cheryl Watson's TUNING Letter, from January 1991 through December 2001. We used Adobe® Acrobat® PDF format, which means you can now, with one search, find and display all the past articles containing a key word, phrase, or APAR. It's a knowledge base of 3500 pages and one hundred and thirty-seven documents at your fingertips, knowledge that Cheryl worked ten years to compile!

You'll also receive the **next six issues** of the TUNING Letter via **email**, each one a searchable **PDF file**, ready to be shared by everyone at your location. And in January, we'll send you the **Annual Updated CD-**

**ROM**, which will include all issues published in 2002 as well as other presentations, speeches and articles.

The electronic TUNING Letter is, like the print edition, offered as a site license, to be distributed to everyone at the receiving physical location. What if you have more than one location? With our electronic edition we offer **MULTI-SITE pricing**. For instance, you may distribute the files to up to four locations for only \$895 per year. Even more locations? Call us for pricing or visit our Web site at <http://www.watsonwalker.com> for all the details.

**Put Cheryl's knowledge at your fingertips today!**

**Yes, we'd like to subscribe!** We would like to receive the new CD-ROM, six future issues of Cheryl Watson's TUNING Letter via email, and the updated CD-ROM in January.

*As soon as payment is received, we will mail you the new CD-ROM containing all back issues of Cheryl Watson's TUNING Letter. Please note that there are no refunds for electronic subscriptions.*

### Single Site:

\$695/year

### Multi-site Pricing:

Send us one CD-ROM and six PDF files to the master address, but allow us to make the files available to:

- 2-4 sites \$895/year
- 5-10 \$1995/year
- 11-50 \$4995/year
- 51-100 \$9995/year
- 101+ \$19,995/year

The EMAIL ADDRESS for delivery of the six future PDF issues is: (please print)

\_\_\_\_\_

Your Name \_\_\_\_\_

Address \_\_\_\_\_

City/State/ZIP \_\_\_\_\_

Telephone \_\_\_\_\_

FAX this form to 941-924-4892  
or **subscribe ONLINE** at [www.watsonwalker.com](http://www.watsonwalker.com)  
or send EMAIL to [admin@watsonwalker.com](mailto:admin@watsonwalker.com)

- We'd also like to order \_\_\_\_\_ extra copies of the CD-ROM: Single-site - \$195 each, Multi-site - \$95 each. (Electronic subscribers only)
- We'd also like to continue receiving a PRINT copy of the TL. Add: \$95, per subscription, \$125 outside of North America.