

Cheryl Watson's TUNING Letter

1998, No. 6

A PRACTICAL JOURNAL OF S/390 TUNING AND MEASUREMENT ADVICE

Inside this issue...

My heart attack is important news, at least from my point of view. See page 2. But I'm feeling really great these days, thanks to modern medicine.

Upgrading a processor is the focus of this issue. How to size, how to understand the difference between speed and capacity, and how to avoid typical problems after an upgrade are all covered starting on page 14.

Java and Component Broker are featured in two articles provided by Glenn Anderson of IBM Education (page 6).

IBM is now recommending that multi-system sites be at OS/390 R5 before freezing systems for Y2K. See WSC Flash 98044 on page 10 for this important item.

A known integrity exposure in ISPF has existed for the five years since ISPF V4, but new customers keep running into the problem. See page 29.

New Web links and important new manuals and books are listed in our S/390 News starting on page 12.

Don't go to OS/390 R5 without checking with your TCP/IP vendors or you could be in serious trouble. See page 28.

Several installations are having serious CPU overhead problems after migrating applications to Y2K compliance. See my article on page 32.

How to set LPAR weights, why fixed storage abends occur, and common WLM questions are addressed in our Q&A.

Benchmark results of the performance APARs for COBOL in an LE environment using CICS are highlighted in the CICS Performance Tips by **Bob Archambeault**.

This Issue:

A Note From Cheryl & Tom	.2
Management Issues #203
Class Schedule4
S/390 News	
GRS Star5
MXG & CA-MICS5
CMOS & Compression5
Y2K6
Java6
Component Broker7
LE9
WSC Flashes10
The Net11
Pubs & APARs12
Focus: Processor Upgrades	
Sizing Processors14
Migration Issues23
Upgrade Problems26
Cheryl's Updates	
Most Common Q&As28
TCP/IP28
ISPF Exposure29
WLM Update30
Y2K32
Q&A	
LPAR Weights34
WLM Questions35
Fixed Storage Shortage	..35
CICS Performance Tips	
COBOL, CICS & LE36
Cheryl's List	
#1739



A Note from Cheryl

If you were a subscriber to Cheryl's List on December 10, you are aware that we were robbed at gunpoint in our home (see Cheryl's List #17 on page 39). At the time I wrote that, we were in the process of recovering, and I was experiencing shortness of breath. Tom, diagnosed with a torn rotator cuff, was just beginning physical therapy. The day after I wrote that Cheryl's List, I found myself in the emergency room having trouble breathing. They diagnosed 'congestive heart failure,' where fluid builds up in the lungs because the heart can't keep up. Part of my heart wasn't moving. It seems that my regular doctor had failed to deduce from my complaints that I had suffered a heart attack the day of the robbery. I lost 50% percent of my heart's pumping capacity as a result of the stress of the robbery. (Just stress – they didn't find clogged arteries.) After two stays in the hospital I became stabilized via various blood thinners and blood pressure medicines. My heart remains damaged, but the medicine helps to alleviate its workload. I'm feeling fine, but tire a lot more quickly now. I'll certainly be taking things easier. Apparently I don't handle stress well!

This hasn't changed any of our key plans, though. I'm still writing the TUNING Letter and supporting our software product. (**Note - Since this was published, Cheryl has made an almost complete recovery. She's working from home and feeling great.)

I will still be at SHARE in San Francisco February 21-26 - see <<http://www.share.org>>. We won't have a booth, but I'll be giving two sessions:

- 2594 (Monday 3pm) - Cheryl's WLM Quickstart Policy Update & Recommendations and
- 2543 (Friday 9:30am) - Cheryl's Hot Flashes (this won Best Session at the last conference!)

Many thanks for all the cards and email you sent after the note about the robbery. We haven't had a chance to get caught up and reply, but Tom and I really appreciate your kind words.

...And from Tom

One of my main jobs as Cheryl's partner has always been to help her manage her time and pick priorities – she's an "A type" personality who tries to do too many things too well, and she always puts her customers' welfare before her own. It seems like we've been cutting back on classes, travel, and newsletter issues almost since we started. Now, not only does she deserve to take it easier, she *must*, and I'm going to make sure she does. Thank you all for your notes of support and concern. You form a great community, one we're very proud to belong to.

As you may have read recently in our electronic "Cheryl's List," **WE HAVE MOVED** our offices from downtown Sarasota to a wooded office complex further south. **See page 4.** ■



Please share this section with your technical management. In it we summarize the technical issues raised in this newsletter (or elsewhere) which we feel merit management's attention.

Focus: Upgrading a Processor

I seem to get more questions about processor upgrades than any other topic. Part of this interest is due to the high acceptance of the very fast CMOS processors. Both IBM and Amdahl have seen higher than anticipated shipments of their CMOS machines and HDS is shipping new Skylines all the time. The Focus in this issue is on processor upgrades - how to size them, why you might not see the expected capacity, and how to avoid the most common problems after an upgrade. "Problems after an upgrade? Why would that happen?" you say? It's a matter of expectations versus observation. In three articles starting on page 14, I explain why the speed and capacity of a new processor can be quite hard to predict.

Important Warnings

This issue contains several important warnings for you:

- MVS/SP 4 and ISPF V4 introduced an integrity exposure for library updates that still exists and is hurting some sites. Page 29 describes the problem with the View command and how to control it.
- TCP/IP for OS/390 R5 was re-written and eliminates support for a function used by several ISVs (Independent Software Vendors). If you're planning to go to R5, you must check with your vendors (page 28).
- One of the more important items in this issue is the notice about a WSC Flash W98044 (Positioning for Year 2000). IBM has provided a document that explains why *multi-system* sites that are planning to freeze operating system releases should stabilize on at least OS/390 R5 instead of R2 as many sites are currently doing (page 10).
- You should *never, ever, ever*, place a production coupling facility in an LPAR that shares CPs with other LPARs. Page 28 describes some WSC flashes that help you understand the reasons behind this warning.
- Many installations are seeing an increase of 10% to 30% in their programs after converting to Y2K. I've listed some ways to avoid the overhead on page 32.

Cheryl's Updates

In this section (page 28), I provide updates on some topics that I've previously discussed. In addition to some of the warnings above, I cover the following items. 1) The catalog CPU overhead problem is now documented in a single APAR. 2) IBM seems to have eliminated the discrepancy in COBOL packed decimal performance with their G5 (9672-Rx6) processors. 3) Workload manager's "small consumer" logic has frustrated some people - there's a description and status. 4) Web Server classification rules are incorrect in the WLM manual - there's an update from the WLM developers.

Elsewhere

Our Q&As on page 34 answer three very common questions: how to set LPAR weights; why fixed storage problems occur; when will WLM compatibility mode go away; and what is the real overhead of WLM goal mode. Our S/390 News provides good news on the following: tremendous performance benefits of using GRS Star; how to get a free cross-reference between MXG and CA-MICS files and variables; why some people are seeing overhead in hardware data compression; causes for CPU overhead during Y2K migrations; some feedback on further LE problems; a final solution to the OS/390 R4 JES2 \$ACTIVATE problem (WSC Flash W98017A); some nifty new Web sites; and several new books worth your notice on UNIX, Lotus Domino, and S/390. Many sites were shocked to hear of the CPU increases in CICS after LE implementation. **Bob Archambeault** reports on five new IBM APARs that result in outstanding performance improvements for CICS and LE on page 36. ■

Cheryl's Class Schedule

Advanced OS/390 Performance & Capacity Planning

**May 17-21, 1999
(Last class ever!)
Sarasota, FL
Register early!**

OS/390, Parallel Sysplex & Workload Manager

***No longer taught;
ask us about video tapes
of last course***

**Call us at 800-553-4562
or visit our Web site at
www.watsonwalker.com**

Our New Location

**Watson & Walker, Inc.
2477 Stickney Point Road
Suite 207B
Sarasota, FL 34231**

**Our 800 number remains the same
800-553-4562
but our local numbers have changed
and are as follows:**

**Telephone: 941-924-6565
Fax: 941-924-4892**

*O*ur S/390 News provides an ongoing update of APARs, important announcements, new publications and other news that we think will be of interest.

GRS Star

GRS Star Outperforms Ring

A reader (who prefers to remain anonymous) sent in the following results of his conversion from GRS Ring to GRS Star (requires a coupling facility). The results were pretty impressive:

"Thought you might like to see some numbers from converting a five member GRS ring (with RESMIL 5) to a GRS star using a 9674-C05 processors.

"Running a program that did a STCK, ENQ, STCK, DEQ, STCK on a unique QNAME-RNAME as a batch job I saw the following numbers:

"In a ring: Mean ENQ time 55.1 milliseconds, mean DEQ time 51.8 ms, min ENQ time 38.2 ms, min DEQ time 31.7 ms, max ENQ time 73.3 ms, max DEQ time 112.2 ms.

"In a star: Mean ENQ time .9 milliseconds, mean DEQ time .6 ms, min ENQ time .5 ms, min DEQ time .3 ms, max ENQ time 2.2 ms, max DEQ time .96 ms.

"Average ENQ improvement 59 to 1, average DEQ improvement 90 to 1."

MXG & CA-MICS

Have you ever wanted a cross-reference between CA-MICS and MXG files and variables? It's now available free online as a service by Nicus Software Inc. on their Web site at <<http://www.nicus.com>>. Nicus is well-known for their products and con-

sulting services in IT accounting and chargeback, capacity planning and performance management. The cross-reference is part of a Web section that helps people convert from CA-MICS to MXG and includes sample code and steps for conversion. The cross-reference itself, however, can be used to do a lookup either way (e.g., you have the MXG name and want to know which CA-MICS name it corresponds to).

I think this is a wonderful service that will be valuable to many installations. Don't miss some of the other goodies on their Web site, such as the free downloadable programs. One creates a comma separated variable (CSV) file from any SAS data member. Thank you, Nicus!

CMOS & Compression

Change in Performance

I have received a few complaints from users of the G4 and G5 CMOS processors about higher CPU time than expected for hardware data compression (used by such exploiters as DB2).

My sources in IBM confirmed that there has been a design change that could affect performance if used frequently. Hardware data compression was introduced with the 9021-711 family. It consists of an instruction that compresses or expands data. In the bipolar machines, it was implemented in hardware. With the exception of the G3 (which also implemented this instruction in hardware) a microcode solution was provided to perform this function on all other CMOS systems. The microcode implementation was chosen to optimize single chip performance (i.e. there wasn't room on the chip for the compression instructions as well as more important items).

What some users are seeing is an increase in relative CPU time (mainly for data expansion) when moving from a G3 CMOS to G4 or G5. However,

even though the G4 and G5 have slightly different internal designs, you should see the expected G4 to G5 CPU ratio reflected also in hardware data compression (i.e. about half the CPU time). Denser technologies will allow IBM to implement the hardware solution once again in future CMOS machines starting with the G7 (9672-Rx8). They expect many more sites will be exploiting hardware data compression by then.

Y2K

Setting GMT

Andrew Stewart of Stelco Inc. sent a note about the following:

"We have cloned our production system using SNAPSHOT copy and we are currently certifying our software for Y2K at various dates and times. One key item you may want to mention to people who are future dating partitions that share a single physical CPU with current dated systems:

"You absolutely must specify the GMT offset when setting the clock. On our first GO we simply set the local time, ie.

```
R 00,DATE=2000.001,CLOCK=07.00.00
```

"The local time was accurate, however we had several products that still reported the date and time as 1998 07.00.00 (COBOL, SAS, Insight, Abendaid).

"Then we found APAR **PQ10688** and set the clock as follows:

```
R 00,DATE=2000.001,CLOCK=11.00.00,GMT
```

"In a 4 hour offset (clock00) this would set the local time to 07:00 am and all our products were reporting the correct date (2000)."

Date Formats

Andrew also ran into the following:

"For SMF and RMF and processing records, both IFASMFDP and ERBRMFPP used the same date format prior to Y2K, ie. DATE(yyddd,yyddd)
I guess IBM had two different teams working on Y2K for these products because the date formats have changed for Y2K and they are different.

"For program IFASMFDP use the following date format: ie. Jan 03,2000

```
DATE(yyyyddd,yyyyddd)
```

```
DATE(2000003,2000003)
```

"You cannot use DATE(00003,00003) since it defaults to 1900.

"For program ERBRMFPP use the following date format: ie. Feb 28,2000

```
DATE(mmddyyyy,mmddyyyy)
```

```
DATE(02282000,02282000)
```

"No other format will work for ERBRMFPP. So now we have two different formats for two very closely related products."

Java

The following section was provided by Glenn Anderson (IBM Consulting Instructor) on the IBM Education Technical Tips page. See their site at <<http://www.training.ibm.com/ibmedu/techtip/>> to find other helpful tips like this.

Servlets? Applets? What's The Difference?

One of the uses of Java on OS/390 is for writing back-end programming for your Web server. A lot

of people are doing that today. If you're one of them, you should know about servlets.

What's the difference between a servlet and an applet? When your server loads a little piece of Java code down to the browser, that's an applet. Servlets stay on the server.

Here's a servlet scenario: A Web request comes from a browser and arrives at the Web server that's running on your OS/390 system. If the Web request triggers a Java program to run on OS/390 itself, as part of the Web server environment, that's a Java servlet.

A Java servlet could use the Java function called Java Database Connectivity (JDBC) to access data that's going to be sent back to the browser. The Java code would make use of JDBC to pull some data out of DB2, for example, stick it in a Web page, and send it back to the browser.

There are two immediate benefits to servlets. You're writing in Java, which people want to do. Also, there is a direct support connection inside the Domino Go Web Server to support these Java servlets. That means the servlet is running in the Web server address space, as opposed to creating additional address spaces. That's good for performance.

There is new support for servlets in the latest Web server in OS/390. The servlet support is actually built into the Domino Go Web Server, Release 5, which is brand new. It's available on OS/390 Version 2.5, which was released in March of 1998.

If you're running Java code on OS/390, or are thinking about it, you should also know about Component Broker, an object-oriented environment coming to OS/390 in 1999. You don't need Component Broker to use servlets, but Component Broker will take care of many nuts and bolts services and also let you do things with Java that you can't do at all today.

For more about Component Broker and Java capabilities on OS/390 platforms, check the next item on Component Broker.

Find Out More About Java On OS/390

There's a brand new Redbook on Java on OS/390, **Integrating Java with Existing Data and Applications on OS/390 (SG24-5142)**. Look at the end of the next article for classes relating to Java.

Component Broker

*This is another article by **Glenn Anderson** which can be found at the same Web site mentioned in the previous article.*

Coming To Your Local S/390: Full Object Environments

You can write Object-Oriented (OO) programs under OS/390 today, using Java and C++. But the mechanics of communicating with other OO programs -- or with legacy applications -- are up to you to deal with. That's due to change.

Component Broker, a sophisticated OO deployment environment, is slated for availability under OS/390 in 1999.

Two Reasons For Mainframe OO Applications

Why worry about OO applications on a mainframe? Two big reasons:

1. The OO design paradigm promotes a better business model in your applications. OO design also offers productivity gains through better reuse of code.
2. As Java programs and other OO applications grow into key business strategies, the scalability

and robustness of mainframe platforms can greatly extend their horizons.

Why A Better Business Model?

OO design principles promote a better business model.

In OO, you're always thinking of the data and the actions you take on the data as a single entity. That mirrors the way business processes actually work.

In an inventory, for instance, you add items to inventory, remove items from inventory or count the inventory. Those are business process. The inventory itself is not a business process.

That's one benefit. Productivity gains through code reuse is another. OO design principles promote the re-use of code and designs.

So, as an organization trying to be more productive, and trying to design better applications for the business, OO is a design model you may want to choose.

Separate Business Logic and Data Retrieval

Component Broker will be a new functionality in OS/390. You have all the advantages of OS/390 in a full-fledged OO applications environment. You can scale your OO applications to much larger sizes, and you've got Workload Manager to manage the environment.

In the OS/390 Component Broker environment, you'll be able to separate your application into business objects and data objects. Business objects handle the business logic, and data objects provide the connectivity between the business object and your existing operational data, whether it's in DB2, IMS, CICS, Oracle, or other sources.

The connectivity of reusable data objects running under Component Broker will shield the business

object developer from worrying about how to access the data.

Built-In Middleware Supports Connectivity

Built into Component Broker is a standard piece of middleware that can connect across a network to other object environments. That functionality is called the Object Request Broker (ORB). Component Broker's ORB will comply with Common Object Request Broker Architecture (CORBA) standards.

When you're writing data objects, you can take advantage of Component Broker's ORB functionality. The data object will call an ORB method that will know how to go across a network and retrieve data for you. The ORB knows about things like opening a TCP/IP socket and building a Sockets connection across the Internet to TCP/IP on another platform.

There's also a naming service, so applications can call for objects and the ORBs will find them on other platforms.

The first release of OS/390 Component Broker, in 1999, will support C++ and Java programs. Other OO languages will follow. Component Broker is currently available on AIX and Windows NT platforms.

Find Out More About OO And Component Broker

If you're looking at future OO projects, Component Broker will be the way to go. To find out more about Component Broker on OS/390, schedule a visit to the OS/390 Expo and Performance Conference in Orlando October 18-25, 1999. There will be lots of OO topics on the agenda. There will be sessions on the basics of object technology, and on Component Broker, along with a three-part Java programming workshop.

If you're looking at a current Java or C++ project on OS/390, check out the IBM course, An Introduction

to Object Technology for Technical Professionals (N1606). And watch for a course coming next year on OS/390 Component Broker.

LE

A reader sent in the following item after reading our 1998, No. 4 issue on LE:

"Cheryl, first let me say "GREAT job on the issue with LE/370"!! I didn't see anything about this surprise in your TUNING letter so I am letting you know about another "little" gotcha.

"We have Abendaid (from Compuware) without the XLS feature installed on our processors. When we introduced LE370 there was a significant change in where the Abendaid formatted dump was placed. LE/370 traps the abend condition and calls Abendaid as a condition handling exit. At this time MVS has not opened SYSUDUMP, it doesn't even know an ABEND is in progress.

"Abendaid dynamically allocates a new DD called ABENDAID to hold the formatted dump. It will model SYSUDUMP or CEEDUMP if they exist but only to a certain point, at least in JES3. Our production jobs make extensive use of an IBM product called RMDS (Report Management and Distribution System) to retain dump information so programmers have a consistent place to find these things when they get called. They use an OUTPUT statement to route to RMDS and an OUTPUT parameter on the SYSUDUMP DD to refer back to the statement.

"In JES3 Abendaid does not have access to information needed to model this parameter or statement for the dynamically allocated ABENDAID DD. In this case the default DD parameter of SYSOUT=* is used.

"We have many jobs that do not include a MSGCLASS= parameter on the jobcard. This is

because of the meticulous way they have of routing all DD statements to where they belong. Shortly after making LE the default for COBOL runtime requests there was an ABEND. The programmer got called and there was no dump information in RMDS. Since SYSOUT=* was used and no MSGCLASS= parameter was present and the OUTPUT statement did not have a DEFAULT=YES the dump routed to a printer in the computer room due to the system wide default. Since the dump was "lost" the programmer had to drive in and resolve the problem without the assistance of the dump information. The next day the four inch stack of blue bar was delivered to his desk. He was not very happy about the whole thing as you might expect. The first reaction was, well, just add the ABENDAID DD card to your jobs. This is easier said than done. We have several thousand JOBS each with multiple steps so this is not a trivial undertaking. We are probably going to recommend a MSGCLASS=x and a DEFAULT=YES on the SYSUDUMP OUTPUT statement.

"This should pick up the "surprise" DD names. CEEDUMP will be another one that gets dynamically allocated when we change from TERMTHDACT(QUIET).

"This is not a highly technical problem or slick approach but is something that can cause ill will between the systems folks and the applications folks, especially at 2:00 in the morning."

In a later email, the reader mentioned the following:

"While this problem was between ABENDAID and LE, your readers might be interested in APAR **PQ22684**. It will provide support specification of a SYSOUT destination if a CEEDUMP DD card is not provided in the JCL."

WSC Flashes

This section provides a summary of IBM's Washington Systems Center flashes. You can get the full text of these from your marketing rep or from a new Web site at <<http://www.ibm.com/support/techdocs>> and click on "Flashes." This new Web site also contains pointers to white papers and FAQs. An example of one of the white papers is one called "Customizing JES2 with Dynamic Tables and Control Block Extensions," by John Hutchinson of WSC. This describes new facilities to customize JES2 that are now available in APAR **OW32032** and will be available in OS/390 R7.

W98042 - CS for OS/390 V2R6 TCP/IP Migration Tips

This 15-page document provides migration recommendations for anyone migrating to Communications Server for OS/390 V2R6 TCP/IP. This release of TCP/IP is the second release of the rewrite that was provided in V2R5, and is only available on OS/390 R6. If you're responsible for installing TCP/IP on OS/390 R6, you'll definitely want to get this well-written document.

W98044 - OS/390 Positioning for Year 2000

This flash documents an online manual that helps customers determine the appropriate software levels for Y2K freezes. The URL is <<http://www.s390.ibm.com/stories/year2000/coexist.html>>.

The main target of this document is customers who have multiple systems that must co-exist or sites who believe they will be able to stay on OS/390 R2, R3, or R4 until 2000, then easily migrate to the then current OS/390 release. Without getting into par-

ticulars here (you REALLY need to read the paper), it's important to note one point that is mentioned in the document and I know is not commonly acknowledged or understood:

"IBM recommends that customers with a multisystem complex or Parallel Sysplex configuration be on OS/390 Release 5 or higher by January 1, 2000."

Since OS/390 R5 can no longer be ordered, you may need to migrate to R6 within the next year.

IBM recommends that customers with a multisystem complex or Parallel Sysplex configuration be on OS/390 Release 5 or higher by January 1, 2000

W98045 - Potential Data Integrity Problem

This flash describes the potential for data integrity problems during DFSMSdss full volume copy or dump operations. It indicates the necessity for applying APAR **PN89166** (9/96).

W98046 - 9672-R06 Performance: Dynamic Dispatch Default Set to Enabled

The default setting for dynamic dispatch is set incorrectly for the 9672-R06 and should be manually changed. Information on when and how to change this is included in the flash.

W98047 - TCPWINDOWSIZE for ADSM Windows NT 4.0

This flash provides recommended sizes for the TCPwindow size option for ADSM backup using Windows NT.

W98048 - DFSMSHsm: DASD Data Sets with Special Expiration Dates Expired

Starting on January 1 (99001 in Julian format), I started getting calls from people who found lots of

deleted data sets because of the use of 99000 for expiration dates. Most of the time, these came from tape management systems that used the 99000 as a special flag. The purpose of this flash is to indicate to DFSMS users that if such tape data sets get migrated to DASD, they will also be deleted from DASD. The flash also mentions two APARs that are important to install: **OW37011** and **OW37046**. I hope you weren't a site who got burned on this.

W98017A - JES2 \$ACTIVATE Warning for OS/390 R4 & R5

This flash actually came out earlier in 1998, but was updated in November. I've previously notified subscribers about the problem with JES2 \$ACTIVATE on OS/390 R4 & R5. The APARs to correct the problem are now available and identified in this flash: **OW31341**, **OW29813**, and **OW31452**. APAR **OW32920** can be used to UNACTIVATE until all correcting APARs have been applied. Also refer to APAR **II10760** for additional information.

The Net

Link of the Week

We've changed our Web page to refer to these links as simply "Cheryl's Links." I won't be updating these again until I get a little more caught up. The last two additions to our links are:

Free On-line Dictionary of Computing (FOLDOC)
<http://wombat.doc.ic.ac.uk>

User Web Pages

<http://www.watsonwalker.com/link981113.html>

(There are several pointers here to real users' home pages that provide a wealth of information. One of the important sites is Chuck Hopf's home page at <http://www.chopf.com>. Select his 'Techo-Gibberish' and look for an excellent paper on moving to goal mode, including some SAS code to help you with the conversion.)

Acronym Finder

Tom, my partner and Web surfer extraordinaire, found a nifty site that has over 69,000 acronyms at <http://www.mtnnds.com/af/fr-top.asp>. Did you know that MVS can stand for Multiple Virtual System, Multiple Virtual Storage, and Multivision? There are seven definitions for SRM including System Resources Manager and Solid-Propellant Rocket Motor. It's not a BIG mainframe site since there are only two definitions for WLM: 1) Buoy Tender, Coastal (USCGC) and 2) Wiring List. But I found it pretty interesting anyway.

There is a further list of acronym related databases at http://www.yahoo.com/reference/Acronyms_and_Abbreviations/.

New Search Engine

Tom's also found a nifty new search engine that we both like a lot better than any of the others. The links are put in importance order and makes searching much easier. The site is at <http://www.google.com>. It doesn't index as many sites as the older search engines, but it should soon.

Google Inc. was founded in 1998 by Sergey Brin and Larry Page to make it easier to find high-quality information on the web. The company is based on three years of research in web search and data mining done by the founders in the Stanford University Computer Science Department.

Emoticons & Chat Jargon

If you do a lot of email and want some additional emoticons (e.g. ;-p means a tongue sticking out) or want a definition for chat jargon (e.g. YGWYPF means "you get what you pay for"), try out this site at <http://www.currents.net/resources/dictionary>. It also has other high-tech definitions.

Listserver Reference

Phil Sidler of Airborne Express found a great site for finding discussion lists. At <http://www.lsoft>.

[com/lists/listref.html](http://www.ibm.com/lists/listref.html)> you can search for lists by list name or topic. Search for "IBM" and you find 43 different mailing lists. It's not as comprehensive as it thinks it is, since I couldn't find two of my list-servers, but it's got a bunch!

Tom likes one called LISZT at <<http://www.liszt.com>>.

Net Locations Mentioned in This Issue

Cheryl's List Archives

<http://www.watsonwalker.com/archives.html>

IBM Education Technical Tips

<http://www.training.ibm.com/ibmedu/techtip>

IBM Redbooks

<http://www.redbooks.ibm.com>

IBM SmartBatch

<http://www.s390.ibm.com/products/smartbatch>

IBM TCP/IP Stack Conversion Cookbook

http://www.software.ibm.com/enetwork/commserver/about/api/api_csos390.html

IBM WLM

<http://www.s390.ibm.com/wlm>

IBM WSC Flashes

<http://www.ibm.com/support/techdocs>

IBM Y2K Freezes Document

<http://www.s390.ibm.com/stories/year2000/coexist.html>

IBM Y2K Web Site

<http://www.software.ibm.com/year2000>

Morgan Kaufmann Publishers, Inc

<http://www.mkp.com>

Nicus

<http://www.nicus.com>

SHARE

<http://www.share.org>

Publications & APARs

Reader's Recommendations

While talking to **Gilberto Rodriguez** of *Interactive Systems* during a class I was teaching, he men-

tioned several books and articles that he found extremely useful or interesting. I haven't had time to read all of them yet, but thought I'd pass on his suggestions. Many thanks, Gilberto.

"The article I mentioned to you about S360-370 architecture history is: **Case Study: IBM's system/360-370 Architecture**, David Gifford and Alfred Spector, *Communications of the ACM*, Vol 30, no. 4 (April 1987) pages 291-307. Although this article is presented as a case study, it really is a very interesting interview with Andris Padegs and Richard Case, two of the people responsible for the design and evolution of the IBM 360/370 architecture.

"I found this special article on parallel sysplex on the **IBM Systems Journal**, Vol. 36, No. 2, 1997 - S/390 Parallel Sysplex Cluster. It contains the design overview of different components of parallel sysplex.

"Also the **IBM Journal of Research and Development**, Vol 41, No 4/5, 1997 - IBM S/390 G3 and G4 contains a technical description of G3 and G4 hardware and S/390 microprocessor design. There is also an article on the Floating point execution unit.

"UNIX books:

"**UNIX for the Mainframer**, David B. Horvath, Prentice-Hall 1998, ISBN 0-13-632837-7. It gives many examples and analogies between UNIX, MVS, TSO etc.

"**The Design of the UNIX Operating System**, Maurice J. Bach, Prentice-Hall 1984, ISBN 0-13-201799-7. This provides details on how the KERNEL works."

Computer Books

Morgan Kaufmann Publishers, Inc. (415-392-2665) publishes several books on computer architecture. While little of the information is specifically oriented for mainframes, there are many titles that

cross all computer boundaries. Their Web site is at <<http://www.mkp.com>>.

One pretty interesting book is: **The Benchmark Handbook: For Database and Transaction Processing Systems**, Second Edition, edited by **Jim Gray**, Digital Equipment Corporation. While it's oriented to UNIX servers, the book also explains how the current benchmarks are created and administered, along with recommendations on creating your own benchmark.

Lotus Domino Redbooks

Two news redbooks are available to help with Lotus Domino. The first is **Lotus Domino for S/390 Performance Tuning & Capacity Planning (SG24-5149)**. Also look for **Enterprise Integration with Domino for S/390 (SG24-5150)** which explains how to connect Lotus Domino to a DB2 database.

Items Found Elsewhere in This Issue

Manuals

This list includes all manuals mentioned in this issue with the exception of those noted above:

GC28-1761 - MVS Planning: Workload Management
SC31-8643 - Lotus Domino Go Webserver 4.6.1 Webmaster's Guide
SC31-8691 - Lotus Domino Go Webserver 5.0 Webmaster's Guide
SG24-5142 - Integrating Java with Existing Data and Applications on OS/390
SG24-5168 - Capacity Planning for CICS Web-Enabled Applications on OS/390

APARs

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

Catalog - II10752 (pg 28)
CICS & COBOL - see LE below
DFSMSdss - PN89166 (pg 10)
DFSMSHsm - OW37011, OW37046 (pg 10)

JES2 - \$ACTIVATE - OW31341, OW29813, OW31452, OW32920, II10760 (pg 11); new function - OW32032 (pg 10)
LE - PQ22684 (pg 9), PQ14883, PQ16794, PQ14888, PQ16844, PQ17931 (CICS & COBOL, pgs 36-38)
WLM - OW25827 (pg 30), OW31890, OW31894 (pg 31), OW32140 (pg 35)
WSC Flashes - 9608 (COBOL, pg 28), 9609, 9609, 9723, 9731, 9731 (LPARs, pg 28), 98017A, 98042, 98044, 98044-98048 (pg 10)
Y2K - PQ10688 (pg 6) ■

A Cheap Hex Calculator!

Sometimes you need a calculator that's a bit more advanced than the one that comes with Windows 95. Look no further than the Calculator itself. In one quick click, this seemingly average applet gets smart.

Open the Calculator (if it isn't already) by selecting Start, Programs, Accessories, Calculator. Not much to look at, eh? But now select View, Scientific, and watch it grow! To switch back to basics, select View, Standard.

Focus: Processor Upgrades

One of our most popular articles was "Sizing CMOS Processors" published in our Nov/Dec 1994 TUNING Letter. I refer more people to this article than any other because it's so applicable and it's an issue in every site. It's no longer just a CMOS issue. Since it's more important today than before, I've updated the entire article. The original article concentrated on moving to slower CPs (physical processors) that were introduced with the original CMOS line, but now that people are moving to faster CPs (150 MIPS or so), the other portions of that article are of more importance.

This focus is composed of three articles, all related, each with a slightly different focus. The first deals with how to size processors; the second addresses additional migration issues, such as the difference between speed and capacity; and the third is a summary of problems to avoid after a processor upgrade.

Sizing Processors

The question I'm most frequently asked these days is "How do I size a new processor?" This article should help you in that process.

Is This Applicable?

This article is useful for anyone considering a change in processor size or in the number of processors. It's also applicable to anyone planning to move a workload from one machine to another. The comments in this article apply to moving from an old bipolar to a slower CMOS machine, to a faster CMOS machine, to a Skyline machine, or moving between any of the above (e.g. Skyline to fast CMOS or vice versa).

Finding Candidates

Before you consider upgrading your processor, you'll want to determine a list of candidates based on available options.

Determine Additional Capacity Needed

But before you can even make a list of candidates, you will need to know how much additional capacity is needed. You may need additional capacity for the following reasons:

1. New or Additional Work

This process is too complex to address in this article, but I'll refer you to our 1998, No. 3 issue with its focus on capacity planning. I will assume that you've determined this value. This is normally expressed in terms of "an x-percent increase in peak period capacity."

2. Software Upgrades

Often a move like this involves an upgrade in the operating system version and corresponding software versions. With each upgrade, there is a resulting increase needed in both CPU and storage. For example, to move from MVS/XA SP 2 to MVS/ESA SP 5, you'll need almost 14 MBs of additional central storage plus and additional 83K per address space (for 200 address spaces, that's over 18 MBs).

See Figure 1 for the estimated changes in storage and CPU. This move also results in an ITR (Internal Throughput Rate, a measure of CPU usage) change of between -13% and +12.4%. That is, there's a possibility that your workloads could require 13% more CPU after the upgrade. Of course, they could take less, but it's safer to plan for the worst case and be thrilled if it doesn't occur.

Figure 1 - Storage and ITR Changes in MVS & OS/390 Releases

From	To	Storage Change	ITR Delta
MVS/370	SP 2.1.0	+1.5M + 50K/AS	-7.8% to +7.5%
SP 2.2.0	SP 3.1.0e	+1.5M + 50K/AS	-3% to +12%
SP 3.1.3	SP 4.1.0	+3.6M to +4.6M	-1.6% to +1.1%
SP 4.1.0	SP 4.2.0	+1.2M + 16K/TSO AS	-.8% to 1.3%
SP 4.2.0	SP 4.2.2	no change	no change
SP 4.2.2	SP 4.3.0	+1MB +1MB(TSO) +16K/AS	TSO: -2.1%
SP 4.3.0	SP 5.1.0	+4.5MB + 1K/AS	Batch: -2% CICS V3: 0% CICS V4.1: -4% TSO: -5.5% to -2% IMS: -5%
SP 5.1.0	SP 5.2.0	+1.2 MB + .5K/AS	Batch: -0.6% to 0.3% TSO: -2.5% to -2.1%
SP 5.2.0	SP 5.2.2	No expected change	No expected change
SP 5.2.2	OS/390 R1	See WSC flash 9613 - Virtual storage and auxiliary storage increases expected due to additional products	TSO: 2.91 to 3.75% DB2: 0.02% to 2.87% Batch, CICS, and IMS: Equivalence
OS/390 R1	OS390 R2	+1MB	TSO: -1.19% to -0.8% DB2: -1.7% to 0.29% Batch, CICS, and IMS: Equivalence
OS/390 R2	OS/390 R3	+2MB + 200 bytes/AS	TSO: -2.5% DB2: -0.95% to -0.06% Batch: -1% CICS and IMS: Equivalence
OS/390 R3	OS/390 R4	+1MB	TSO: -1.5% DB2, Batch, CICS, and IMS: -0.6 to equivalence
OS/390 R4	OS/390 R5	none	TSO: -0.03% to -0.4% 0 - rest
OS/390 R5	OS/390 R6	none	Unknown; probably equivalent

Note: A negative ITR value translates to a decrease in internal throughput and an increase in CPU usage of approximately the percentage reported. A positive ITR value translates to an increase in internal throughput. All these are cumulative changes and must be added if you are moving past multiple releases.

Thus, when sizing for an existing workload, the primary thing you need to do is to ensure that you've provided additional storage, and also enough additional CPU capacity to make up for any ITR degradation that may occur. Note that I've only indicated the storage and CPU changes for MVS and not for new versions of other software applications. You may also want additional storage to take advantage of the many new data-in-memory facilities such as hiper-spaces, data spaces, and VLF, provided in the newer MVS versions.

3. Latent Demand

If you are upgrading because you are currently out of capacity, you must account for the latent demand that exists on the system. Latent demand is work that is not currently running on the system but may appear if the configuration changes. Latent demand tends to be two types of work:

- Work that wants to run at a certain time of day but is prevented from running by higher priority work. A good example of this are

long batch jobs that only run at night on a constrained system, but can run during the day after an upgrade. The daytime processing is heavier, but the nighttime processing is reduced. The work simply *moves*.

Other work of this nature can be seen as 'in and ready' or 'swapped out and ready'. In and ready users are ready to execute on the machine, but are not dispatched due to the volume of higher dispatched work. Out and ready users are artificially kept out of the system due to other activity. In either case, when the higher dispatch work is run on a bigger machine, this 'ready' work will start executing sooner.

- Work that is *created* because of a faster CPU. This might be because people can type faster on a faster processor doing data entry type of work. Or it might be because they change the way they work. I remember changing batch turnaround response time at one site from 30 minutes to 5 minutes. We saw an increase of 400% in batch processing because people started using the computer to do more debugging than before. Another type of created work is that which occurs because a new technique is found to work. At another site, processing improved so much in a claims processing unit after a processor upgrade that everyone started using CICS browse for lookups rather than typing out full names. The number of transactions remained the same, but the resources used doubled.

If, for example, you are running a 200-MIPS machine at 100% and you upgrade to a 300-MIPS machine, you might expect the new machine to run at 67% busy with the same work. This will never happen! It will run at 75% or more simply because of latent demand. A simple rule of thumb is to allow a 10%-20% in-

Figure 2 - List of Candidate Machines

MODEL	MSU	MIPS	CPs	MIPS/CP
HDS Sky 527	114	654.4	5	130.9
Amdahl GS885	111	650.0	8	81.3
Amdahl GS877	110	648.4	7	92.6
IBM 9672-Y56	109	642.0	5	128.4
Amdahl GS7Y5	109	640.8	11	58.3
IBM 9672-R66	109	639.5	6	106.6
HDS Pilot 68S	109	639.5	6	106.6
HDS Sky 625	106	620.0	6	103.3
Amdahl GS868	103	604.0	6	100.7
IBM 9021-711	11	63.0	1	63.0
IBM 9021-9X2	78	484.5	10	48.4

crease in CPU utilization (during peak periods) for latent demand. The article on page 27 also addresses latent demand.

Determining Candidates

Look up your current machine in our CPU Chart and find three items: Avg MIPS, # of CPs, and MIPS per CP. Adjust the average MIPS by your needed capacity. For this article, let's assume you currently have an IBM 9021-9X2 and want an additional 30% capacity (for added capacity, software upgrades and latent demand). We list that machine as having 484.5 MIPS, 10 CPs, and 48.4 MIPS per CP (Figure 2). It's also in processor group 80 with a rating of 78 MSUs (Millions of Service Units used for PSLC pricing). A 30% increase over 484.5 MIPS means that you'll want about 630 MIPS. Since the average MIPS is only an estimate, you'll want to investigate machines in the 600 to 660 MIPS range.

The next step is to come up with a list of possible candidate machines. The list of Processors on page 24-25 of the CPU Chart contains all the machines listed in descending processor group, MSUs, and MIPS. From this chart, you can see that the biggest processor group 80 machine is one at 530 MIPS, so you'll need to go into IMLC pricing. (If you don't have our CPU Chart in front of you, trust me on this one.) Now it's a matter of finding possible machines. Figure 2 shows the starting candidates I would select.

Of course, you might want to eliminate some of them due to the vendor, age, or availability. If you are using processor group for software pricing, you might also eliminate any candidates that push you to a higher software group.

For this example, all candidate machines have faster CPs. One has more CPs (11) and the rest have less, as few as 5. Part of your analysis will be to determine how your workload will behave in each of these configurations. Your work will run very differently on an Amdahl 11-way at 58.3 MIPS than on an HDS Skyline 5-way at 130.9 MIPS or an IBM 9672-Y56 at 128.4 MIPS. The following sections should help you understand the difference.

If you are currently planning a processor upgrade, I suggest that you go through this same process by looking up your current machine in our CPU Chart and finding both the MIPS estimate and the number of CPUs. You can either create a list of candidates or, if you've already decided on a specific machine, collect the same data for the machine that you're planning to move the workload to.

Please note that I am going to use the average machine MIPS for my analysis in this article, even though I would strongly advise you to take this analysis one step further by looking at the anticipated performance by workload using the MIPS by Workload column of the CPU Chart.

You'll next want to size your workloads and see if they'll fit on the new machine. I want to discuss some different scenarios and list the considerations for each. To understand which scenario applies to you, you need to understand both the speed and capacity of your current machines. 'Speed' refers to how much work a single CP can process, while 'capacity' refers to how much work can be processed by all of the CPs on a machine. A multi-CP machine can have two speeds, a 'uni-speed' and an 'mp-speed'. The uni-speed is the base speed of a single CP and is normally quoted as the speed for a one-way machine in a series. As an example, an IBM 9021-711 machine has one CP with a speed of 63.0 MIPS (Figure 3). That speed CP is used to build the 10-way 9021-9X2. When all ten CPs are trying to process work at the same time, there is some inter-

Figure 3 - CPU Chart Extract

Machine	#of CPs	Avg MIPS	MIPS/CPU	MP
9021-711	1	63.0	63.0	1.00
9021-822	2	119.1	59.5	0.95
9021-832	3	173.3	57.8	0.92
9021-942	4	224.9	56.2	0.89
9021-952	5	274.7	54.9	0.87
9021-962	6	322.6	53.8	0.85
9021-972	7	366.7	52.4	0.83
9021-982	8	408.2	51.0	0.81
9021-9X2 - 9-way	9	446.4	49.6	0.79
9021-9X2	10	484.5	48.4	0.77
Skyline-115	1	124.0	124.0	1.00
Skyline-215	2	236.0	118.0	0.95
Skyline-315	3	342.0	114.0	0.92
Skyline-415	4	430.0	107.5	0.87
Skyline-225	2	236.0	118.0	0.95
Skyline-325	3	342.0	114.0	0.92
Skyline-425	4	438.0	109.5	0.88
Skyline-525	5	532.0	106.4	0.86

ference and overhead called the 'mp-effect'. The same work will take longer to process. The capacity of the 9021-9X2 is listed as 484.5 MIPS for all processors. Dividing the capacity by the number of CPs gives us the effective speed or mp-speed for that machine, or 48.4 in the case of the 9X2. Since most people think in terms of MIPS, I'll use MIPS for comparison, but you can also use the SU/sec rate as provided in our CPU Chart.

There are different considerations depending on whether you're moving work to a faster processor (CP), to a machine with faster but fewer CPs, or to a slower processor. Here are three different situations to consider:

- 1) Moving to faster CPs
- 2) Moving to fewer but faster CPs
- 3) Moving to slower CPs

1) Moving To Faster CPUs

This case implies that you are moving to faster CPs with the same number of CPs or more.

This is the most common scenario today, since all of the newer processors are faster than ever before. Additionally, many installations that have the older 308x, 438x, and smaller 309x machines are forced to move to newer machines in order to support Y2K changes.

Another reason that many sites are considering this option is the elimination of support for MVS/XA. Sites that don't want to lose the software support must move to OS/390, which doesn't run on some of these older machines.

This is the easiest move of all, and seldom causes any problems except those listed in the Processor Upgrade article on page 26. The biggest problem with faster CPs is simply the increased cost of the software, so be sure to determine how expensive the upgrade will be for both hardware and software before you commit to the upgrade.

If you are moving to a significantly larger capacity machine than you've ever run (over 500 MIPS or so), be sure to read the article on large machines on page 24 for some additional problems.

2) Moving To Fewer But Faster CPs

If you are moving to fewer CPs, with each CP being faster than the older processors and providing more capacity, you might think that there is nothing to worry about. Actually, there are several things to worry about. First, take into consideration all of the items I mentioned in the previous section about moving to faster CPs.

Then you need to look at your workloads. Some workloads really run better with more CPs instead of faster CPs. Take the example of moving from a 3081/G (2 CPs) to a 9672 R11 (one CP). The MIPS for the two-way 3081/G are 10.6 or 5.3 MIPS per CP, while the MIPS for the 9672 are 13.8. This move represents a 30% increase in CPU power, so you should be okay.

But this is moving from a two-way to a one-way that's faster. If all the applications are small applications that don't want to dominate the CPU, there will be no problem. If one of the applications is CPU-intensive and would take more than half of the

3081/G if it could, then you might have problems when you reduce the number of CPs. If the application is high priority, such as CICS, and it had been limited on that 5.3 MIPS CP, it would take more CPU on the 9672 than it had on the 3081. The other applications could suffer.

This phenomenon occurs all of the time, such as moving from a 3090/600J (6-way of 114 MIPS total) to a 9021/821 (2-way of 117.6 MIPS total). In many installations, that type of upgrade won't work. All it would take in this last example is to have two high priority applications that are CPU hogs, and all other work would suffer. The six-way would tend to limit these two applications to no more than 19 MIPS each ($114 / 6$), but the two-way would allow them to gobble up 58.8 MIPS ($117.6 / 2$) if they weren't constrained.

Many installations have been hurt during upgrades like this. One installation combined two machines with a total of twelve processors to a single machine with 30% more capacity but only five processors. It was a disaster! They happened to have four very large online systems that dominated the new machine. On the twelve processors, the online systems were restricted by the speed of a single CPU. They thought the machine would last for a year and it was out of capacity by 10am on the first Monday morning!

In order to see if an application is likely to take more CPU, look to see how much it's using on the current machine. If a single workload (especially a high priority workload) is using a full CP's worth of processor, then it will most likely take more CPU resources on a faster processor. Maybe that's what you want, but you'd better plan for it. A quick way to see how much a workload is taking is to look at an RMF or CMF report during the peak periods for that workload.

The example in Figure 4 shows a sample RMF report by domain. You can start with the domain report and move to the performance group report to obtain more detail about the highest activity domains. "AVG" and "TCB+SRB%" are the two fields of importance for sizing CPUs. (APPL% is used instead of TCB+SRB% in RMF/CMF releases

Figure 4 - RMF Workload Activity Report

WORKLOAD ACTIVITY													PAGE 29
MVS/ESA		SYSTEM ID SYSA		DATE 09/09/94		INTERVAL 14.59.979							
SP4.3.0		RPT VERSION 4.3.0		TIME 12.44.00		IPS = IEAOPT00							
OPT = IEAOPT00		SUMMARY BY DOMAIN		SERVICE DEFINITION		COEFFICIENTS		SU/SEC=1883.2					
ICS = IEAICS00				IOC = 5.0		CPU = 10.0		SRB = 10.0		MSO = 0.0001			
PGN	PGP	DMN	TIME SLICE GROUP	INTERVAL	SERVICE	AVERAGE ABSORPTION, AVG TRX SERV RATE, TCB+SRB SECONDS, %	PAGE-IN RATES	STORAGE	TRANSACTIONS	AVG TRANS. TIME, STD. DEVIATION HHH.MM.SS.TTT			
ALL	ALL	075	ALL	IOC	206.8K	ABSRPTN 205	SINGLE	0.13	AVERAGE 1,262.8	AVG 65.00	TRX	000.00.00.000	
				CPU	10854K	TRX SERV 205	BLOCK	0.13		MPL 65.00	SD	000.00.00.000	
				MSO	491.5K	TCB 576.3	HSP	0.00	TOTAL 82,080	ENDED 0			
				SRB	468.9K	SRB 24.8	HSP MISS	0.00	CENTRAL 40,321	END/SEC 0.00	QUE	000.00.00.000	
				TOT	12021K	TCB+SRB% 66.7	EXP SNGL	2.96	EXPAND 41,759	#SWAPS 0			
				PER SEC	13,356		EXP BLK	0.00			TOT	000.00.00.000	

after 5.2.) "AVG" is the average number of concurrent address spaces. "TCB+SRB%" or "APPL%" is the percent of a single CP that's needed for the workload. Thus, a value of 250% indicates that the workload needs 2.5 CPs.

Note that the TCB+SRB% is only available since SP 4.2. You can calculate this value from older reports by taking the CPU service units and dividing by the SU/SEC and the CPU SDC. Perform the same calculation for the SRB service units and add the resulting values to get the total CPU time for the interval. Now divide by the length of the interval. As an example, use the values from Figure 4. The calculations are:

$$\text{CPU time} = (10854000 / 1883.2) / 10.0 = 576.36 \text{ seconds}$$

$$\text{SRB time} = (468900 / 1883.2) / 10.0 = 24.9 \text{ seconds}$$

$$\text{Interval} = 14.59.979 \text{ (mm.ss.ttt)} = 899.979 \text{ seconds}$$

$$\text{TCB+SRB\%} = (576.36 + 24.9) / 899.979 = .668 = 66.8\%$$

If a high priority single workload takes over 80% of a single CP, it's an indication to you that it will probably take more CPU on a faster processor.

In Figure 5, I've extracted some data from a few key workloads. Notice that PGN 006, CICS, is currently using 97.4% of a single CPU. Since CICS is primarily a single tasking address space (even in CICS V3, over 85% of the work is still done under a single TCB), I would suspect that this CICS region

needs more CPU than it's getting. If you move it to a faster CP, it will take more resources. If you've reduced the number of CPs, CICS may take more resources than anticipated and the lower priority workloads may suffer. Of course, the CICS users would be happy!

In the same figure, PGN 007 is IMS that's also taking almost an entire CP. The AVG field indicates that there are five address spaces in this performance group. In order to see if any IMS region will take more CPU in the new machine, you'll need to analyze each of the five address spaces (perhaps use reporting performance groups for awhile).

When sizing for fewer CPs, look at your primary, high priority applications to see if they will tend to take more CPU resources if available. If so, be prepared to deal with the situation by somehow restricting their access to the CPU (e.g. by lowering their dispatch priority).

3) Moving To Slower CPs

This is the situation that was more common in 1994 with the small CMOS machines, but you will still see some examples of this today.

It's quite possible that you can go from a one- or two-way CEC to a 6-way CEC, where the CPs are slower (fewer MIPS per CPU). Or you simply might be adding a smaller machine and planning on moving some selected applications to it. This last option

Figure 5 - Extracts from RMF/CMF reports

PGN/ Period	Workload	TCB+SRB%	AVG # of AS	Total % of R61	AVG% of 900 per AS	AVG% of R61 per AS
005	Batch	66.7	65.00	247.0	1.0	3.8
006	CICS	97.4	1.00	360.7	97.4	360.7
007	IMS	98.4	5.00	364.4	19.7	72.9
031	Hi Pri STC	22.5	1.00	83.3	22.5	83.0
002/1	TSO 1st period	41.5	2.75	153.7	15.1	56.0
002/2	TSO 2nd	13.5	1.08	50.0	12.5	46.3
002/3	TSO 3rd	34.2	2.62	126.7	13.1	48.4
002/4	TSO 4th	9.4	.84	34.8	11.2	41.4
002/all	TSO total	98.8	7.30	365.9	13.5	50.1

Data taken from 9021/900, 1883.2 SU/sec, 6-way, 228 MIPS
Comparison is to 9672/R61, 508.1 SU/sec, 6-way, total 62.9 MIPS, each CPU is 27% (508.1/1883.2) of a 9021/900

has been chosen by several companies who are trying to elongate the life of their current mainframe by moving selected applications to a new machine.

As an example of moving to a slower CP, let's assume you have a 9021/711 machine (one CP of 63.0 MIPS). When determining the right size of CMOS machine, you'll want to first add in the additional overhead for upgrading the software versions. Let's assume we're going from MVS/XA SP 2 to MVS/ESA SP 5, and therefore must allow for a 13% CPU increase. Adding 13% to 63.0 MIPS gives us 71.2 MIPS needed on the new machine. Be sure to add in some capacity for latent demand as mentioned earlier. For purposes of illustration, let's assume that you'll want an additional 29 MIPS for latent demand and growth, resulting in a 90 MIPS requirement.

Estimating Total Capacity

First determine the optimum configuration of machine that will provide the total capacity. Look at page 12 of the November 1998 CPU Chart. Let's assume that you're considering the IBM 9672-RB5 rated at 88.8 MIPS. Remember that a MIPS rating may vary by 20% or more depending on your type of workload. See the discussion on page 14 for typical sizing techniques. [Historically, this has been the only analysis needed for sizing new CPUs - you need to have enough in the total capacity of the machine. With smaller CPs, however, you'll need to evaluate whether the applications will fit on the new CPUs.]

As an example, let's look at the 9672-RB5. This is a two-way rated at a total of 88.8 MIPS or 44.4 MIPS per CP. Since the current workload is running on a single CPU rated at 63 MIPS, we're looking at a decrease of 30% in speed of a single CP. That is, a single address space running on the 9021-711 might take one CPU second, but will take 1.3 seconds on the 9672-RB5.

You need to understand what this difference will make to your workloads. And all workloads are different! To understand the effect of a slower CPU, you need to have an understanding of the life of a transaction or address space.

Sizing Batch and TSO Applications

Let's take a very simple example of a batch job. Most batch jobs spend the largest portion of their elapsed time waiting for I/O or the CPU (because they run at a low dispatch priority). Typically a small percentage of the elapsed time is actually using the CPU, with a large percentage of the elapsed time waiting for the I/O or CPU on a busy system. Let's look at three jobs that each take ten minutes elapsed time:

```
Job A |Wait for I/O.....|Use CPU|
Job B |Wait for I/O.....|Use CPU.....|
Job C |Wait for I/O.....|Wait for CPU.....|Use CPU|
```

If you run the jobs on a slower CP, the "Use CPU" portion of the job will be increased. If that's only 10% of the job (as for jobs A & C), then only the

10% of the elapsed time will be increased. In my example of a 30% increase, Job A would take 10.3 minutes due to the slower processor. Job B, that is a CPU-intensive job, will take a much longer time. If the "Use CPU" time was 60% of the elapsed time, the job would take 30% more of the 60% of the time and complete in about 11.8 minutes.

Job C, however, would probably benefit from having more CPs, even though they were slower. If job C is spending 20% of its time waiting to get to a CPU, you might be able to reduce the "Wait for CPU" time by having more CPs and thus reduce the elapsed time. Any workload that has a LOT of address spaces, such as batch and TSO, will often receive better response time if there are more CPs, even though they're slower.

How does this longer elapsed time affect your system? In many ways. Be aware that the effects are multiplied, however, (often exponentially) as the number of MVS images increases. For example, shared DASD delays may be slightly noticeable when you add one image to a shared DASD environment, but they'll be significantly longer when you've added the sixth or seventh image!

Kathy Walsh of IBM's Washington Systems Center has a good set of reminders (in italics below) relating to the MPL (multi-programming level - i.e. amount of currently active work) issues of slower engines (this was presented by Kathy and **Linda August** in 1994, but is still applicable today).

Reminder 1. As work slows down, it will hold onto resources longer (e.g. data sets, storage, CPU, etc.).

Reminder 2. Many sins of the past have been fixed by faster engines:

- Work Packs*
- Working Set storage*
- Tape Drives*
- Initiator scheduling*
- Shared DASD*
- Data in Memory requirements (Hiperspaces)*
- Job scheduling across multiple systems*

If your batch window is fairly tight already, lengthening the elapsed time could cause you to miss your deadlines. When profiling your batch workloads, there are other issues to be analyzed in addition to the CPU time and storage usage. Kathy also included a good checklist of things to consider when profiling your batch workloads:

Check ratio of CPU to Elapsed time. If batch is CPU-intensive, focus more on this CPU-intensive workload.

What fixed resources do jobs need, i.e. tape drives?

Does batch work use a data in memory technique, which forces affinity? Would this batch work then cause the 9672 to define some of its CSTOR as ESTOR? What is the performance cost?

What other resources will either cause increased costs or force affinities to certain processors, i.e. compilers, tools, vectors, crypto?

Does batch use a database manager running under a different TCB? Does the database manager have a single TCB architecture?

Can batch scheduling tool handle multiple systems?

Does the current JES initiator structure/job class definitions support the PTS environment?

What are the impacts of shared I/O resources on throughput?

List factors which cause batch window to be missed today. How much allowance is given to re-runs, will it need to be expanded for slower jobs?

Will the TIME= parameters need to be changed to account for longer CPU times? If it's not changed, you may see unexpected S322 abends.

SmartBatch for OS/390 (previously known as BatchPipes/MVS) is a product from IBM that can help reduce elapsed times for batch jobs. It was specifically designed for a parallel environment. It allows concurrent execution of multiple jobs in order to reduce elapsed time. I mentioned Batch-

Pipes/MVS in the July/August 1993 Share Trip Report. For more information on SmartBatch for OS/390, check out their Web site at <<http://www.s390.ibm.com/products/smartbatch>>.

Sizing Online Systems

The online systems are a slightly different matter. Most online applications are fairly CPU-intensive, and any decrease in the speed of the CP will increase the CPU time and thus the response time of the transactions. On the other hand, can your users tell the difference between an internal response time of .30 seconds and .33 seconds when the end-user response time is 3 seconds? Of course, if the CP is only 1/20th the speed, you might be talking about 2 seconds internal versus .3 seconds!

Since most online systems process multiple transactions concurrently, you'll see little I/O delay for a transaction, but you should be aware of the delay waiting for other transactions to complete (wait to dispatch). On a slower CP, the transaction will not only take more CPU time during execution, but it will have a larger dispatch wait time due to other transactions that are also taking more CPU time.

If you need to reduce the effect of the slower CP speed, you may have to divide the online systems into multiple regions to allow them to truly multi-task. Part of the response time of an online transaction is waiting for other transactions. By dividing the workloads and putting them in different regions and on different CPs, you can reduce the effect of the slower speed. Dividing the workload into multiple regions also improves the parallelism of the transactions that can provide benefits in many different situations.

Again, you'll need to look at your important workloads to see how much of the CPU they need. This process is used whether you're replacing an existing machine or simply moving selected applications to a new machine. Look at Figure 5 again and at PGN 006, the CICS region. We see that this one address space (AVG = 1.0) is currently taking 97.4% of a 9021/900 CPU. By dividing the TCB+SRB% by .27 (the relative size of the 9672), we see that the same address space would take 360.7% of a 9672 (3.6 CPUs).

In order to even move this CICS workload to the 9672, it would need to be separated into at least four address spaces. Of course, any time you increase the number of regions, you'll need to take into account the additional CPU resources needed for sharing (about 15%), as well as additional storage (close to 100% more for each region!).

Summary

In sizing processors, keep in mind the following:

1. Consider the increased capacity needed for newer versions of the operating system and your subsystems, as well as any change in capacity due to your business plans.
2. Initially size the machine by total capacity (e.g. MIPS) in order to ensure that there is enough total capacity for all of your workloads.
3. Look at each important workload. Look at the usage per CP on the current machine in terms of MIPS and percent of the CPU.
 - a. If the workload is taking close to a full CPU and you're moving to a faster CPU, consider that the workload may take more of the total capacity than on the older machine.
 - b. If you're moving to faster engines but with fewer CPUs, try to determine if any of the workloads will tend to dominate the CPU at expense of the other workloads.
 - c. If the workload needs more MIPS per CPU than is available on the new machine, determine if you can tolerate the longer elapsed or response times. For online systems, consider providing multiple address spaces to handle the transaction load, but don't forget to add the needed capacity for the splitting of regions.
4. Measure your applications closely both before and after any upgrade. If one workload is experiencing a performance problem, it might be attributable to another workload that's getting too much of the system.

5. Don't forget to consider the simple upgrade of your current machine. While it's pretty exciting to move to the new technology of CMOS processors, many installations can save time and money by simply extending their current bipolar machine. If an installation needs an extra 50 MIPS of processor power, it's much easier to upgrade to a larger bipolar (unless you're sitting with a ten-way machine of IBM's or a twelve-way of Amdahl's) than it is to move work to multiple CMOS CPUs.
6. Adding another CP to a current machine takes little, if any, effort from the systems support staff. Adding another machine means the addition of new systems volumes and new DASD, a sysgen process, channel configurations, additional hardware for channels and CTCs, a configuration process, time spent on determining workloads for the new machine, shared DASD analysis, new license fees for most software licenses, and changes to reporting programs. Slapping in another CP is a piece of cake in comparison. (Does "piece of cake" translate as "super simple" in other countries as well?)

I'd also like to recommend some CMG papers that address the issue of sizing different numbers of processors:

Hackenberg, Steven R., "More Engines, Faster Engines, and Partial Engines," CMG 93 Proceedings, p. 619 (this material was also published in Enterprise Systems Journal, January 1994)

Krause, Irwin, "The Virtual Processor: Capacity and Performance Analysis in a Partitioned Environment with Fractional Engines," CMG '93 Proceedings, p. 793

Krause, Irwin & Leganza, Gene, "Continuous System Measurement: Self-Referential Analysis for Insight Into Your Computing Environment," CMG 94 Proceedings

Leganza, Gene, "Rating Processors Without Benchmarks: An Introductory Tutorial," CMG 93 Proceedings, p. 19

Migration Issues

For sites which have moved to faster, but fewer CPs, I'd like to address three additional issues that might come up after an upgrade. These also apply to sites moving to a larger machine (i.e. adding CPs). This topic has generated a large amount of email. You should be familiar with the material in the first section of this Focus article before proceeding.

These additional issues are:

- 1) Speed Versus Capacity
- 2) Large Machines
- 3) Increased Expectations

Speed Versus Capacity

A common problem occurs after an upgrade because people tend to misunderstand the difference between a machine's speed and its capacity. Most sites look at the CPU usage of a workload to determine how it is doing. They might look, for example, at the amount of CPU used by a common and well-understood CICS transaction. The CPU usage is really a measure of the speed of the CP and the architecture, such as multi-processing (MP) overhead, that comes into play.

Most hardware performance guarantees (you did get one, didn't you?), on the other hand, are based on capacity, the total deliverable MIPS of a machine. When the apparent speed of a machine appears to differ from the capacity, disagreements often occur between the vendor and the customer. I think the problem lies in the definition and expectation of speed versus capacity.

As my example, I'd like to use a case of moving work from an IBM 9021-9X2 to an HDS Skyline 525. The same situation can be seen moving to a new G5 CMOS, but the numbers for the 9X2 to 525 are much easier to use for explanation.

In the case of the 9X2 versus the Skyline 525, using the values from my November 1998 CPU Chart (as shown in an extract in Figure 3), the two machines have the following characteristics:

9X2 - 484.5 average MIPS capacity; 48.4
MIPS/CPU relative speed as a 10-way; uni-
speed is 63 MIPS

525 - 532.0 average MIPS capacity; 106.4
MIPS/CPU relative speed as a 5-way; uni-speed
is 124 MIPS

Please note that I am going to use the average machine MIPS for my analysis, even though I would strongly advise you to only analyze estimated capacities and speeds by workload (CICS, TSO, etc.). Since I'm doing this as an example of technique only, the difference between using average machine MIPS and MIPS by workload is irrelevant.

The relative speed takes into account the multiprocessing (MP) overhead as shown in the accompanying figure. From a speed point of view, the 525 has 2.0 (124/63) times the uni-speed, but 2.2 times the relative speed (106.4/48.4), and only 1.1 times the capacity (532.0/484.5 or 10% more).

We can see that the Skyline is twice the uni-speed of the 9X2, but provides only 5 CPs instead of 10 CPs. If the MP-factor were not taken into account, the two machines would provide almost the same capacity (630 MIPS [10*63] versus 620 MIPS [5*124]). But the MP-factor does need to be taken into account. This is the additional overhead incurred when work is moved between CPs during execution.

There is between a 3-5% loss in speed and capacity for each CP added. Therefore, the 10th CP on the 9X2 causes all CPs to run at only 77% of their uni-speed (see the MP column in Figure 3) and the 5th CP on the 525 causes all CPs to run at 86% of their uni-speed ($86/77 = 1.1$ or 10% better). This MP-factor is what allows the 525 to provide 10% more capacity on the average than the 9X2. What this tells us is that even though the 525 has only double the speed, it can provide 10% additional capacity because of the MP-factor.

But what happens to specific workloads? Let's take an example of DB2 (or CICS or IMS). These workloads typically run at the highest dispatch priority in the machine and use a lot of cycles. There-

fore, once they start execution on a CP, they don't easily give up control - they sit on a single CP and use it. In effect, they run at uni-speed and don't get hit much by the MP-effect. So what occurs is that the very best that DB2 can do is 2.0 times the uni-speed. You can measure this by looking at the difference in CPU time for known transactions. DB2 (like CICS and IMS) will not see the expected 10% extra capacity.

The lower priority workloads, like batch and TSO third period, will see the improvement because they only getting hit by a 5-way MP-effect on the 525 (86%) instead of the 10-way MP-effect of the 9X2 (77%). If you ran the entire 9X2 and 525 with only DB2, I would expect that you would start to see the 10% improvement in capacity on the 525. As it is, you don't see it because DB2 is not hit hard by the MP-effect (which is where the 525 gets its boost).

What does this mean to you? If you want the improved speed purely for your DB2 workload, then your contracts in the future should be based on speed (and even capacity) of your loved workloads, such as DB2. If you want the additional capacity for lower priority workloads, then you cannot use the observed speed of the DB2 workload to infer the capacity and speed of all workloads.

All of this discussion of the 9X2 compared to the Skyline is not restricted to these machines. It will occur every time work is moved to a machine with faster, but fewer, CPs.

Large Machines

Early customers of Amdahl's 12-way 5995-12670M (530 MIPS) and HDS's 8-way Skylines (780 MIPS) found that it wasn't easy to increase the size of a machine past 500 MIPS. They were trying to run IBM software, which hadn't been designed or tested on machines that large. In most cases, the customers had to partition the machines and run smaller MVS images.

The problems often had to do with built-in limits (number of concurrent address spaces, initiators, DASD drives, channels, below 16MB storage, etc.). As IBM is creating and designing larger machines themselves, the limits are being removed in the

software. Today, OS/390 can sometimes be run in a single image 1000 MIPS machine. Most people don't do this, however, because some limits (like the 16MB constraint) keep popping up. More programs simply require more storage below 16MB and it's a limited resource. If you put too many of these in storage at one time, you'll run out of "below the line" storage. Most installations with over 500 MIPS that I've seen tend to partition them into smaller images.

What most people tend to forget is the additional overhead of the additional work that can be placed in these huge machines. IBM developers have long said that LPAR overhead is really a multi-programming overhead - that is, the overhead of having more address spaces in storage at the same time. Now we're seeing that multi-programming overhead in single image systems running on very large machines.

It's a fact that as more work is added to a system, the currently running work will take more CPU time due to the additional workload. Just think about it for a minute. CPU time will increase when you increase the number and size of tables, queues, and control block chains; when you increase the chance of getting knocked off your CP by work at a higher priority; and when you increase the number of interrupts because of additional address spaces.

I've confirmed this phenomenon many times. You can run a job alone on a machine and it takes 10 minutes of CPU. If you add two more jobs to the machine, the first job might take 10.5 minutes. If you add 200 more jobs to the machine, the first job might take 12 minutes. The CPU time increases as more work is added to the system. The higher increases are felt by the lower priority work on the system, but all will experience some additional overhead.

Now consider the case where you have a 300 MIPS machine that's running at 100% and a 100 MIPS machine that's running at 100%. Will they both fit on a single 400 MIPS machine? Probably, because you'll save some duplication of effort (i.e. only one JES running, only one SMF running, etc.) and you'll

save some storage. That is, you'll get more capacity because you're running less total work.

But what happens to the speed? To simplify my example, let's assume all three machines had the same relative speed of CPU. The relative CPU time for all of the jobs will increase because of the increased work that they each compete with. The higher priority work will get less of the overhead (for the same reasons I explained above), but all will experience some overhead because of the increased tables, queues, etc.

Increased Expectations

Here's a topic bound to cause disagreements in any installation. When you upgrade your processor, you will most often improve the response times and turnaround times for your users. They'll love you!

But what if that increased capacity was bought because of planned increases in workload or new applications? What happens to the user's response time when that new workload comes into the system? It will increase again and the users will hate you!

This is a common scenario and one that I think is avoidable. I personally believe that you should not give all of the new capacity to users if you plan to take it away again. The easiest method is to place the image in an LPAR, give it a weight that is needed to handle today's workload, and cap the LPAR so the users can't get more. That will leave some idle CPU capacity until the new work appears on the system when you can increase the weight and, therefore, the capacity of the image.

The primary argument against this is that you are throwing away CPU time that users could have access to. My experience is that once you "loan" your users that new capacity, you will never be able to take it back again without complaints. By that time, too, they might have changed their behavior and require the new level of resources.

I'll get off my soapbox now, but if I could convince installations to use this technique more often, I know that the number of processor upgrades would be reduced.

The next section discusses some of the things that change during an upgrade.

Upgrade Problems

I get *lots* of emails from people who tell me heart-breaking stories about how things went wrong with their processor upgrades. Here's a list of how to avoid the biggest mistakes when upgrading a processor:

Performance Guarantee

Insist on getting a performance guarantee from the vendor. If you don't and the machine fails to meet your expectations, then you have no recourse. Be careful of the penalty, however, since it is easy for the vendor to simply upgrade your machine if the capacity is less than expected. That is not really good enough. Most sites cannot afford to accept a CPU upgrade because of the increased software cost. I think the only valid penalty clause is one that allows you to return the machine after a period of time and would allow you to bring in a different machine (even a different vendor) without penalty.

While I recommend that you get a performance guarantee, I also realize that it is very difficult to confirm capacity. The guarantees should be related to the performance of your critical workloads. As I mentioned earlier, performance guarantees are normally based on capacity projections, but most people notice what's occurring with a specific workload (and that is normally determined by speed, not capacity).

You can't easily make judgements about the speed of a new machine without really understanding the underlying relationships. In most vendor hardware performance guarantees, the vendor places a restriction by saying that the workload must remain the same. This is almost impossible to do. You are either moving to a faster processor where the work will change characteristics (or the users will change

their behavior), or you're changing the number of CPs which will impact the number of concurrent users.

Limiting Parameters

Don't let parameters limit the capacity of a new machine. Most sites forget to update their parameters when upgrading and will artificially limit the amount of new work on the machine. The most common parameters that are forgotten are:

You simply cannot use CPU busy as an indicator of either capacity or speed after a move

1. CNSTR in IEAIPSxx in compat mode (limits the number of swapped in users).
2. RCCCPUT in IEAOPTxx in compat mode (can limit the number of dispatchable address spaces).
3. Resource group maximum in goal mode (limits the amount of CPU used by some service classes).
4. VTAM LUs for TSO users (limits the number of TSO users).
5. Storage limits or artificial limits, such as storage isolation values in compat mode and transaction class limits (can limit the number of transactions).

These (and many more) parameters can be found in my two-part series in our 1998, No. 1 and No. 2 issues.

Software Costs

Complete an analysis of the software costs associated with any hardware upgrade. In many cases, the software costs far exceed the hardware costs, yet many sites still forget to do a complete analysis until they've already committed to a specific processor. You should also consider some of the benefits of two processors connected with a coupling facility to qualify for both PSLC discounts on software used on all machines and a reduction in costs for software that only needs to be run on one machine. I'm

aware of at least four sites which found that two machines plus the coupling facility resulted in a significant reduction in total outlay for the upgrade.

The increase in software costs is causing many installations to find replacements for products that are marketed by a select group of intractable software gougers (pardon me, I meant to say "companies"). However you choose to deal with the software issue, it cannot be ignored.

Latent Demand

As I recommended on page 15, make it your business to understand latent demand in order to understand why a processor upgrade seems to take more capacity. One of the most common questions I get relates to a condition where the CPU busy is much higher than expected after a processor upgrade.

As an example, a 400 MIPS machine is currently running at 100% busy during peak period. The site upgrades to a 500 MIPS machine where they expect to be running at 80% busy with the same workload. I can guarantee this will never happen. I would expect, instead, for the new machine to be running at 100% during peak period too. The reason is latent demand. This is work that is currently being limited by the current capacity and finally has a chance to run (for example, a faster processor allows data entry clerks to enter data faster and process more transactions). It could also be work that is normally run during off-peak hours and finally has a chance to run during peak period (test batch jobs, for example).

In the first example, you can actually measure the increase in work on the system by looking at the increase in transactions. In the second example, the total activity for the day may not increase, but the peak period will increase. For this reason, any analysis of capacity should include both a peak period transaction rate and a daily total transaction rate. You can either use transaction rates from your subsystem, such as CICS, or you can use I/O rates as an indicator of the amount of work you've processed.

You simply cannot use CPU busy as an indicator of either capacity or speed after a move.

Summary

I've listed the primary reasons for complaints after a processor upgrade. You may also run into some minor problems, such as increased expectations as I mentioned on page 25 or a misunderstanding of LPAR configuration. If you understand these issues, however, I think you'll be in a much better situation to evaluate what might be happening to your system after an upgrade. ■

Cheryl's Updates

This section provides updates on topics that I've written about before. I've discussed some of the items in presentations at conferences or in my classes, but have never expanded the topics in as much depth as I've wanted to. I hope you'll find these updates extremely useful.

Most Common Q&As

Here's a list of the most common questions I receive, and my answers. Some of these have appeared in previous issues, but the APARs have often been updated in the intervening years.

Catalog CPU Overhead

Starting with MVS/SP 5 (and some say SP 4.3), users have been seeing a significant increase in the amount of CPU time consumed by the Catalog address space, CAS. There are many reasons for this, and an APAR, **II10752**, has been issued to help define the known problems. It also describes how to report apparent increases to IBM. Please do this if you have some concerns about it. I receive a lot of email from people who have not considered reporting the problem to IBM, who, understandably, won't put resources into unreported problems.

COBOL & CMOS Overhead

Shortly after the first CMOS machines began to be used, several customers found that some COBOL programs were taking up to five times the amount of expected CPU time when run on an IBM or HDS CMOS machine. This was the result of a redesign of the CPU chip that gave an advantage to floating point instructions at the expense of packed decimal instructions. WSC Flash **9608** describes the problem from an IBM point of view. You can find additional comments from me on our Web site at the "Cheryl's List" archives, <<http://www.watsonwalker.com/archives.html>>.

Starting with the IBM G5 series (9672-Rx6), this disparity in performance for packed decimal has been eliminated.

But in all cases, poorly written COBOL programs should be changed to use indexes rather than subscripts, regardless of being run on a CMOS or bipolar processor. If you look at the benchmarks in the flash, you can see that the performance is greatly improved by using indexes. (I was teaching this technique to students in the 70s - it's not a new thing!)

LPARs and Coupling Facilities

You should **NEVER, EVER, EVER**, place a production coupling facility in an LPAR that shares CPs with other LPARs when doing data sharing. IBM continually recommends against it, and sites continue to do it to save money. I really can't stress the importance of this recommendation enough. To help you understand why this can be such a problem, IBM's Washington Systems Center has issued several flashes. Look for WSC flashes **9609** (CF Reporting Enhancements to RMF), **9609** (LPAR Performance in a Parallel Sysplex), **9723** (Parallel Sysplex Performance XCF Performance Considerations), **9731** (Dynamic CF Dispatching), and **9731** (XCF Performance Considerations). Note that even though these flashes are almost three years old, they are still applicable to today's machines.

TCP/IP

Vendor Incompatibilities

The Communications Server for OS/390 R5 (VTAM and TCP/IP) in March of 1998 included a complete rewrite of the TCP/IP stack. The performance improvements are significant (in some instances, 15 times better).

There is one major detraction from this great new performance. The new stack no longer supports IUCV and VMCF APIs. Vendors who were members of IBM's free Partners in Development program were notified in 1996 about this change, but several ISVs have never heard of the change or haven't been able to ship R5 compliant programs.

IBM indicates that "customers moving to CS for OS/390 V2R5 do not have a 'work around' option supported or endorsed by IBM. V2R5 compliant ISV products are essential to successful customer upgrades to the new release." So before considering implementation of CS R5 or R6, be sure to contact all of your TCP/IP vendors and confirm their compliance.

An excellent guide for conversion is available on the web, TCP/IP Stack Conversion Cookbook, at: http://www.software.ibm.com/enetwork/comms/erver/about/api/api_csos390.html.

ISPF Exposure

In ISPF V4, you may look at a file in one of three modes: Edit, Browse, and View.

Edit Mode

In Edit mode, you have access to full edit facilities, such as edit macros and line commands (e.g. Exclude or Copy). When done editing, you may SAVE (replace) the member you are editing. You are protected because only one person at a time may have the member opened for Edit mode.

Browse Mode

In Browse, you do not have access to any edit facilities; you CANNOT change or alter the member; you can't save the member; but several people may have it open in Browse mode at the same time.

View Mode - Careful!

The View facility, which became available with ISPF V4, combines features of both edit and browse. The primary ISPF menu will now show

View and Edit as primary options (instead of Browse and Edit). Selecting View from the primary panel actually takes you to a screen that supports both Browse and View. A slash (/) to the left of the words 'Browse Mode' on the screen automatically puts you in Browse mode for this and all future sessions until the slash is removed.

So what's different between View and Browse and why would you use one over the other? Multiple users may open the same member in View mode; you may use the full set of edit facilities as in Edit mode; and you can change and alter an in-storage copy of the member. Even though you cannot do a SAVE to replace the member, you CAN do a REPLACE, which does the same thing. This leads to an integrity exposure, as I'll explain later. Like in Edit, you can also do a Create to create a new member.

View is often used to allow a "temporary" change to a member of a protected library. You might want to make a temporary change to some JCL, submit it, then ignore the changes. Or you might want to take a protected member, use some macros (like the SAMPLIB Cut and Paste) to create a new file before creating a new member with it. The temptation is great to use View, but its very power can make it dangerous.

The REPlace option allows an integrity exposure. Multiple people may have the member open in View mode; multiple people could be altering it; and multiple people can be doing a REPLACE. So if Jack opens FILEA in View mode and Jill does the same thing, they won't know about the other person. Jack sees a change that MUST be made and does it; so does Jill. Jack tries to do a SAVE (thinking he's in Edit mode) but can't when he's in View mode, so he does a REPLACE. Jill does the same thing. Jill's copy overlays Jack's and Jack is clueless that his changes are lost.

Although I had mentioned this integrity exposure in our Nov/Dec 1995 TUNING Letter issue on ISPF V4, many people are still unaware of the problem. All most people saw was that the Browse option on the ISPF menu changed to View. Some sites have

eliminated use of View in the data center altogether because they've been burned by it.

Disabling View

As mentioned earlier, an individual can specify whether they want their default option to be Browse or View by simply entering a '/' in the 'Browse Mode' option at the bottom right of the View screen. Many people disable View to eliminate confusion. One user's comment: "I got sick and tired of editing files only to find out I was in VIEW instead of EDIT. In my VIEW it stinks."

Doug Nadel, author of TASID and ISPF developer, provided the following information:

"View can be disabled via the ISPF configuration table but there is no direct provision for just disabling replace. Here is an indirect method:

"ISPF for OS/390 V2.5 (ISPF 4.5) contains two new edit macro entry points. The ISPF configuration table can be used to define a site-wide edit macro. A variable called ZUSERMAC in the profile or shared pool can be used to define a user (session wide) initial macro, and then there are the usual initial macros. The site-wide one runs first, the user one runs second and the regular one third. The site-wide one could check to see if the session is VIEW and then issue an ISREDIT DEFINE REPLACE DISABLED. The kicker is that there is not a reliable way to tell if you are in VIEW mode. That will come eventually, but a modification of the ISREDDEx panels could be made to save that information in the profile or shared (preferable) pool."

View Storage Usage

One more major difference with View is that it initially brought the entire member into storage (instead of using temporary files like Edit).

A field called EDITSTOR can be set in the ISPF configuration table to define the maximum storage allowance for new edit or view sessions. By setting the field to non-zero, edit/view will keep track of the amount of storage being used for the edit or view session when the data set is being loaded into the editor. No checks are made during the session.

If an edit or view session exceeds the limit set in the configuration table, the edit/view session will be terminated and a browse session will start instead.

This can impact the system performance significantly if you attempt to view or edit a very large data set. For example, if you were to go into View mode on an online data set that's a print of an SVC dump (which is an EXTREMELY large file), you could cause paging on the system to increase significantly. If you know you are going to look at a very large file, use Browse instead of View or set the EDITSTOR field so that it's done automatically. Some data centers saw an average increase of 2 MB in their TSO swap sets when they went to ISPF V4 using View as the default.

WLM Update

Small Consumer

I've written about the concept of the "small consumer" several times in the past, but would like to summarize it and provide the latest update on it.

WLM developers learned early on that WLM could be more effective if it could just get some of the "little" stuff out of the system quickly without much management. In order to do this, it identifies service classes who are "small consumers" (as defined by internal algorithms based on the CPU usage of the address spaces in a service class) and bumps their dispatch priority (DP) to x'FD', just below SYSSTC. In the original implementation, the DP of small consumer was x'FE' and SYSSTC was x'FD', but that was changed with APAR **OW25827**. These small consumers, therefore, get a big boost of CPU access, where they will typically take only a fraction of the CPU before they go into a wait. If they are truly small CPU consumers, the work will get out of the system quickly without impacting anyone else.

Why can this be a problem? Just think of what might happen when a single address space classified as a small consumer decides to use LOTS of CPU (goes into a loop; does a big table search, etc.)? You

will end up with a low importance address space running at a higher DP than CICS, DB2, and other loved ones for 10 seconds or more. Ten seconds is the length of time before WLM re-adjusts the DP. Depending on your environment, systems with 10 or 12 CPs will not see this occur as often as systems with one or two CPs. When the system is running at maximum utilization, however, all systems could see this problem arise. This is especially true when you have filled your CPs with high importance work and the low importance work starts to suffer. WLM might try to bump the low work to small consumer priority.

I get lots of email from people who go to goal mode and are concerned when they find low importance work running at a higher DP than high importance work. For most sites, the small consumer logic works fine - even with some negative reactions from people who look at the DP. For other sites, especially those with only a few CPs, low consumers have been known to severely impact higher priority work.

As one example, **Eric Donohoe** and **Carole Miller**, from The Royal Bank of Scotland plc, reported a problem with small consumers running higher than CICS. CICS response time went to 20 seconds! They have applied the APARs mentioned a little later in this article and are awaiting the results.

In some cases, you can work around the problem. Some solutions to keeping small consumers out of the way:

1. Put notoriously erratic or high CPU users into discretionary (they aren't moved to high DP).
2. Group work into large enough service classes. That is, if you put each CICS region in its own service class, WLM might think a test service class is a small consumer until someone actually started using it heavily. Or if you put several active regions in the same service class, then their activity and need for resources will prevent small consumers from being dispatched at a high DP.

Here's what happened to **Chuck Hopf** (be sure to see Chuck's paper on WLM as mentioned on page 11):

"I had defined three groupings of STC. Hi, warm and low relatively speaking. The problem is that the sum of the HI was often less than 1% and the sum of the low was often less than 1%. Where I got into trouble was when one of the low (importance 4 velocity of 30 or so) woke up after a period of hibernation and like the grizzly bear after hibernation proceeded to eat everything in site (at SC priority.) NETVIEW is one such and there are others that can lie dormant but suddenly become voracious."

Chuck found that by assigning more work to HI and LOW, they stopped being categorized as small consumers.

3. Apply the maintenance listed below to help the small consumer logic.

Gail Whistance of WLM development provided the following status of small consumer:

"The two recent APARs involving high importance work being momentarily delayed by lower importance work are **OW31890** and **OW31894**. Only the latter involves small consumer logic. Both were code errors in WLM's projection logic. They are both closed now and there has been no feedback positive or negative from customers, although around twenty customers were on the interested parties list. These fixes received additional levels of testing on native systems with various types of workloads, such as CICS, batch, and TSO.

"We have not had any field situations involving small consumer logic since OW31894 was opened, although there has been a lot of concern expressed about this concept at conferences. Something customers can do to help WLM be more responsive during periods of high utilization is to review their service classes, especially the high importance ones, and combine any that are handling similar kinds of work and can be assigned the same goals."

Classifying Scalable Web Server Transactions

Gail Whistance of WLM development suggested that our readers would be interested in the following update about classifying Webserver transactions for goal mode. Here is how to set classification rules for the IWEB subsystem:

"If you recently installed IBM's Web server and are classifying your scalable Web server transactions, be careful what documentation you use. The book **MVS Planning: Workload Management (GC28-1761)** has inaccurate descriptions for some of the classification qualifiers for Web server transactions. Use the corrected descriptions below:

"For more information on Web server transaction classification, in particular how to use the Web server configuration file, see **Lotus Domino Go Webserver 4.6.1 Webmaster's Guide (SC31-8643)** or **Lotus Domino Go Webserver 5.0 Webmaster's Guide (SC31-8691)**.

"For a practical example of using the classification qualifiers, see a new redbook, entitled **Capacity Planning for CICS Web-Enabled Applications on OS/390 (SG24-5168)**. This redbook will be available in hardcopy within a month, and is viewable/downloadable now from the web at: <http://www.redbooks.ibm.com> by doing a search on 'SG24-5168'.

"Classification Qualifiers for the IWEB Subsystem Type:

Userid - The Web server's userid (not the original requestor's userid). It is implemented this way because the requestor's userid is not available to the Web server at the time the transaction is classified. This probably has limited usefulness, since the userid is the same for most transactions. The default userid for the Web server is WEBSRV.

Subsystem Instance - This qualifier is not supported by the Web server at this time. It may be supported in the future at which time it will use the subsystem name used in the application environment definition, since this is unique for

each instance of the Web server. For now, don't use this one.

Subsystem Parameter - The actual format is:

0-7	Subsystem name (same as what is planned for Subsystem Instance above)
8	blank
9-23	Source IP address
24	blank
25-39	Target IP address
40	Blank
41-46	Target port

Transaction Class - This is probably the most useful qualifier because of its flexibility. Transaction class is the arbitrary class name you specify in the APPLENV directive in the Web server configuration file. You can use the filtering function in the Web server to assign transactions to transaction classes based on the requested URL. Then in turn, the transaction classes can be assigned unique service classes via the WLM policy using the transaction class qualifier.

Transaction Name - The method name, for example, GET, HEAD, POST, PUT and DELETE."

Many thanks for this update, Gail!

Y2K

Y2K Overhead

As I mentioned in the issue on Language Environment (LE) in our 1998, No. 4 issue, many sites are seeing an increase of 10% to 30% in their programs after converting to Y2K. I'm listing the most common reasons here to help sites determine if this might be a capacity exposure for them.

1. Inexperienced COBOL programmers are making changes. Many of the Y2K programmers have been trained for two weeks in COBOL and placed in an installation to make the Y2K changes. They don't know about things like effi-

cient data conversion or using indexes instead of subscripts. A friend of mine who had retired and hadn't looked at COBOL in twenty years is now contracting out at \$150 an hour updating COBOL programs. It's kind of scary!

2. Most of the date routines are being changed from static (just a branch) to dynamic (a program needs to be loaded into storage). Dynamic processing can add a significant amount of time if not done properly. A simple thing like obtaining today's date for every transaction can cause the CPU time of a job to double when going from static to dynamic.
3. As I mentioned in our 1998, No. 4 issue on LE, you might see a 5 to 30% increase in CPU time if LE and the programs aren't properly tuned.
4. New software often takes more CPU time. New compilers or new releases of subsystems (e.g. CICS, IMS, DB2) can sometimes take more CPU time than the older releases. Again, with tuning, the new releases can often take less CPU time, but you need to put in some amount of effort.
5. Y2K migration resources are seriously underestimated. In almost every installation I've interviewed, they've found that they've underestimated their migration resource requirements by anywhere from 30% to 1000%.

What can you do about these problems? There are several things that can help you avoid most of the problems that other sites are seeing.

1. Measure applications before and after their conversion. See what the impact was. Almost every site runs a duplicate run with a production run to ensure that the output is identical. Just be sure to measure the change in CPU time for each job and determine whether more tuning is needed.
2. Before hiring any consultant, be sure to confirm their skill level. This also goes for situations where you send a group of programs out of house to be modified. Part of your contract

could state what performance degradation you are willing to accept.

3. Tune your applications while you're changing them. Use good coding techniques, such as those recommended in the COBOL Tuning white papers on the COBOL Web site. Tune LE as recommended in our 1998, No. 4 issue. Train programmers on good coding techniques. After all, you expect these programs to keep running for at least two years.
4. During testing, use an application monitor (such as Strobe from Programart or InTune from Boole & Babbage) to analyze where programs are spending their time.
5. Track Y2K projected resource usage with actual resource usage to determine if there will be enough resources during the heaviest part of the conversion.

What Release Works for Y2K?

IBM has announced that OS/390 R1.2 was the first Y2K compliant release, BUT...

1. There are LOTS of year-end APARs that must be applied to R1.2 to resolve items that were not discovered before its release.
2. This just applied to MVS, not necessarily to other products, such as CICS, IMS, COBOL, DB2, etc.
3. New APARs are being released all the time.
4. For multisystem sites, IBM is now recommending that OS/390 R5 is the oldest release you should be on if you're going to freeze updates until 2000. See the WSC Flash **98044** mentioned on page 10.

Be sure to check the PSP bucket YEAR2000 for the latest APARs and check out IBM's Y2K Web site for the status of IBM and non-IBM products, <<http://www.software.ibm.com/year2000>>. ■

Q & A

In this section, we provide common questions and their answers.

LPAR Weights

Gilbert S. Reamy of Philip Morris USA and Galen Leung from Sun Life Assurance of Canada both had the same question about how to set weights for LPARs.

The following TUNING Letter issues provide a lot of information on LPARs:

February 1991, FOCUS: PR/SM
March 1991, Reducing I/O Elongation
May 1991, Amdahl's MDF
December 1991, Dedicated and Shared LPARs
July 1992, Defining a Sysprog LPAR
Sept/Oct 1992, PR/SM Overhead
May/June 1993, PR/SM Changes (SP 4.3)
July/Aug 1993, FOCUS: LPAR Update (EMIF, MDF, SP 4.3)
1997, No. 4, LSPRs and You
1998, No. 1 & 2, Configuration Changes

I last wrote about PR/SM in depth in 1992 and I don't think I mentioned anything specifically about setting the weights for PR/SM, so here are my comments in a nutshell:

Remember that the weight is used to calculate the percent of ALL shared CPs. If you are sharing all CPs, then the weight represents the percent of the box.

The weights of all active LPARs are added together, then each LPAR's weight is divided by the total to determine the percent of the box that the LPAR may get if everybody wants the entire machine. It's a LOT easier if you assign weights in such a way that they normally add up to 100, since it makes it easier

to analyze the results. For example, you might set the weights for three LPARs at 70, 20, and 10. You could also set them at 280, 80, and 40 and get the same results. But in the first case, it's easier to simply look at the total percent busy and see how you're doing (e.g. if an LPAR is taking 80% of the box and its weight is 70, then it's getting 10% more of the CPU than you had intended).

Capping used to take a lot more overhead, but that's been fixed for years and isn't a concern unless you have a much older bipolar machine. Capping tells the system to not allow the LPAR to get more than its percentage as calculated by the weight. If you don't cap and the other LPARs don't need CPU, one busy LPAR can get more than its share.

What's more important than anything is to assign the minimum number of logical CPs (LPs) to each LPAR. That is, if you have a 10-way box, don't assign all 10 LPs to each LPAR; just assign the minimum number to handle the maximum load. That is, if you think one of the LPARs might take up to 70% of a 10-way, then assign it 7 LPs.

The number of LPs assigned defines the maximum amount of the system an LPAR may get and the weight defines the minimum amount of the system an LPAR may get if it wants it. The two are really related. Let's take the 10-way example again. If I set a weight of 70 and assign 7 LPs, then I'm saying that it can have at least 70% of the system (weight), but not over 70% of the system (LPs). A weight of 60 and assigning 7 LPs says that it can have at least 60%, but will never exceed 70%. A weight of 80 and assigning 7 LPs is silly, because it can never get over 70%.

As a side note, inactive LPARs add no overhead (at least it's too small to really measure).

WLM Questions

Here are two commonly asked questions that I get. They happened to be answered by WLM developers on their home site, so I'm including the answers from that site. It's at <<http://www.s390.ibm.com/wlm>> (select FAQs).

Q: I have not yet switched to goal mode. How long will IBM support compatibility mode?

A: "IBM has not yet announced plans to eliminate compatibility mode. It will not happen before the year 2000 for obvious reasons. After that, it depends on what proportion of customers have migrated to goal mode and on our development priorities. IBM does in fact intend to eliminate compatibility mode eventually, but ample lead time will be given so customers have a chance to migrate to goal mode at a reasonable pace."

Q: What is the overhead of goal mode over compatibility mode?

A: "There is rampant misinformation circulating on the overhead of goal mode. The direct additional CPU processing in goal mode is typically around 1/2 to 2%. The exact cost for your environment will depend on multiple factors such as number of address spaces, amount of I/O, and volume of CICS and IMS transactions to sample. Some users, in fact, have seen a reduction in CPU time in goal mode because of better systems management over what they had specified in compatibility mode.

"The one area to watch out for, though, is the MAXTASKS setting for your CICS regions. If MAXTASKS is set arbitrarily high, you incur unnecessary sampling overhead because a sampling control block is created for each potential task. Therefore, limit MAXTASKS to your high water mark plus a buffer. APAR **OW32140** (closed in 12/98) will reduce the overhead of sampling the CICS regions if you are using velocity goals for the CICS regions and not managing the CICS transactions to response time goals." This APAR is applicable to all releases from SP 5.2 to OS/390 R7.

Fixed Storage Shortage

Scott Stanley of Lehman Brothers wrote: "Over the last few weeks we have been experiencing an alarming number of "IRA400E" conditions. In the last two days we lost our Development system twice, and we're starting to see this condition on our production systems as well, so I think you can understand my trepidation regarding the situation."

Even though I discussed this issue in our 1997, No. 6 issue (page 31), I still get several questions regarding this problem.

There are two primary reasons for IRA400E. The first is normally due to a misunderstanding about the RSU= parameter in IEASYSxx. We discussed this in our 1997, No. 6 TUNING Letter, page 31. If this value is not zero, a problem often occurs after a memory upgrade. If you don't have that issue, just send your fax number to <admin@watsonwalker.com> and we'll fax you the article.

The second reason is that some job is taking LOTS of fixed storage and your RMF Mon II (or CMF) data should show that. Collect data from both swapped in and swapped out users, looking for a high number of fixed frames.

[Editor's Note: The reader took a look at RMF frames and found that TCAM (which was not being used) was taking 2 Mb of central storage below the line. It was stopped and the IRA400E messages disappeared!] ■

CICS Performance Tips

Bob Archambeault of R.A. Solutions Int'l Inc. has been providing these wonderful CICS performance tips to our readers for several years and they continue to be a powerful addition to the TUNING Letter. Our many thanks to Bob.

Bob is currently providing migration information and marketing support as an IBM business partner. Bob recommends that readers contact him BEFORE they order the next version of CICS to get the latest information on migration problems and performance information. There is no cost as long as they contact him before they order the next version. He can help them make the right decision about which version to use. That's an offer that's hard to refuse!

--

Bob Archambeault
R. A. Solutions Int'l. Inc.
6209 Pentridge Ct, Raleigh, NC. 27614
919-845-8774
bobarch@ibm.net

(All of the material in this article remains copyright © R.A. Solutions International, Inc.)

COBOL Performance using LE with CICS

I have been receiving many comments on my CICS and LE articles that appeared in TUNING Letters 1998, No. 1 and reprinted in 1998, No. 4. Those articles showed a dramatic increase in the CPU times when using the EXEC CICS LINK commands in COBOL programs after the implementation of LE. I suggested using dynamic calls instead of EXEC CICS LINK commands if a program is linked-to multiple times. I realize that it is not always possible to re-code existing applications to use dynamic calls. Sometimes, the applications may be vendor programs and you do not have access to the

source. Other times, different languages may be involved in the programs used in a single transaction. Since dynamic calls are not always supported between different languages, this is not necessarily the solution for everybody. This article is a follow-up to those articles and provides some very good news.


Many thanks to the folks in CICS, LE and COBOL development for responding so quickly with a solution to the increase in CPU times when using EXEC CICS LINK commands. Five APARs listed below are now available and the results have been outstanding. I've summarized the information contained in those APARs in the descriptions that follow.

1. LE APAR **PQ14883** implements a new option called 'Skip Exit DSA Processing at termination'. In a CICS region with LE, each time a program issues an EXEC CICS LINK command, a new LE execution environment, or rununit, is created. This involves a fairly significant amount of resources and storage when compared with previous language runtimes, such as COBOL II. COBOL II used to extend the rununit for a CICS LINK command instead of creating a new one. LE provides better program isolation and allows for the tuning of runtime options on a program by program basis with its architecture. There are several optimizations that are implemented here in the area of initialization and termination of LE enclaves and threads in order to reduce the overhead of LE for the CICS LINK command.

The creation and deletion of an LE 'rununit' (enclave/thread) under CICS is improved with the following specific changes:

- Faster initialization of the LE Enclave Block (EDB)
- Small optimizations in stack initialization
- Bypass of the global CEEBXITA if it is

Figure 6 - COBOL Link Performance



the default

- Faster termination by allowing languages to skip exit DSA processing
 - Several other optimizations in the initialization and termination path
2. COBOL APAR **PQ16794** makes Language Environment COBOL run-time changes to enable the 'Skip Exit DSA Processing at termination' support provided in APAR PQ14883.
 3. LE APAR **PQ14888** allows the LE/CICS Run-Unit Work Area (RUWA) to be reused across EXEC CICS LINK commands for ALL31(ON) applications, if the size is large enough, eliminating the CICS GETMAINS for this area on every LINK command. When this occurs, LE is able to reuse and reset some control structures instead of re-initializing them. This combination reduces the pathlength of a CICS LINK by approximately 15% for ALL31(ON) run units. Several other LE/CICS interface improvements reduce the pathlength of the initial run unit for transaction initiation.

The LE/CICS language interface has been improved to minimize CPU usage for the creation/deletion of 'rununits' under CICS. The following are some of the optimizations:

- Elimination of CICS call to LE for Rununit

Begin Invocation. LE's Rununit Initialization now calls Begin Invocation directly. The CICS trace entry for Begin Invocation is removed.

- Faster Rununit initialization due to the merge of Rununit Work Areas (RWAs) into a single CICS GETMAIN at task initiation time. CICS optionally keeps track of the HWM of RWAs for a task with its first execution and uses that information on subsequent invocations to minimize RWA getmains.

- Several optimizations in LE that allow for certain resources to be reused/initialized faster.

4. CICS APAR **PQ16844** changes CICS so that, with the PTF of APAR PQ14888 (mentioned above) applied to LE, RUWAs are reused on repeated invocations of LE-conforming applications. A combined call to LE for rununit_initialization and rununit_begin_invocation has also been introduced to reduce the pathlength of the processing involved in starting an LE program. A new System Initialization Table (SIT) parameter, RUWAPOOL=(YES | NO) has been added to allow you to implement this new function.

Tests were done by IBM for all supported releases of LE using a CICS benchmark. This benchmark had 100 tasks, each issuing 5000 LINK commands and writing the results to temporary storage. Figure 6 is a graph of the CPU times for a CICS system using LE 1.8. Every release of LE showed similar results.

You may have noticed in Figure 6 that the CPU times for ALL31(OFF) applications is higher than ALL31(ON) applications. If the LE option ALL31(OFF) is specified, stack storage (C/PLI

automatic variables, COBOL LOCAL-STORAGE, and COBOL library routine work areas) are allocated below the 16 megabyte line. In addition, AMODE switching across calls to LE routines is required at a performance cost.

CICS regions often run ALL31(OFF) since the AMODE requirements of all the programs run in that region is usually not known. Thus, AMODE 31 programs are needlessly run with the default ALL31(OFF) unless explicitly overridden by CEEUOPT. This results in a performance and below line storage penalty being paid.

5. LE APAR **PQ17931** will allow an installation to specify ALL31(ON) to be used as a CICS default in a mixed environment of AMODE31 and AMODE24 programs (that do not have CEEUOPT overrides). LE will determine which programs are AMODE24 and force those to ALL31(OFF) as an installation default override. If CEEUOPT or exits have an ALL31 setting, they will continue to be used in normal order of precedence regardless of the program being AMODE24 or not. Note if an installation uses dynamic calls from an AMODE31 to an AMODE24 program then they must still use an installation default of ALL31(OFF) or use a specific override using CEEUOPT or exits. AMODE24 autodetection will not work for a dynamically called program.

I highly recommend that you apply all five APARs and specify RUWAPool=YES in the SIT for best performance. We have also found improvements by specifying the following parameters for the CEECOPT table for CICS regions. It is important to specify storage sizes in bytes rather than in K-bytes to allow for the CICS storage check zones which are added to best utilize the 4K page structure.

ALL31=((ON),OVR),
ANYHEAP=((4080,4080,ANYWHERE,FREE),OVR),
BELOWHEAP=((4080,4080,FREE),OVR),
HEAP=((4080,4808,ANYWHERE,FREE,4080,4080),OVR),
HEAPCHK=((OFF,1,0),OVR),
LIBSTACK=((512,1008,FREE),OVR),
RPTOPTS=((OFF),OVR),

RPTSTG=((OFF),OVR),
STACK=((4080,4080,ANYWHERE,KEEP),OVR),
STORAGE=((00,NONE,NONE,0K),OVR), * see note below
TRMTHDACT=((MSG),OVR)

* Note: the third STORAGE parameter will need to be 00 if running PL/I or C. The fourth STORAGE parameter is not used by CICS, but if specified will cause an unnecessary getmain for the amount of storage specified below the 16 megabyte line.

Conclusion: LE is a good thing. Initial performance problems are being resolved quickly. Don't be afraid to use it after the above suggestions are considered. ■

Cheryl Watson's **TUNING** Letter

1998 Number 6 (Vol. 8, No. 6)
Published 6 times a year
by Watson & Walker, Inc.

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

SUBSCRIPTION RATES: \$465 per year (6 issues).
Outside North America, \$515 per year. 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., 2477 Stickney Point Road, Suite 207B, Sarasota, Florida 34231, USA. Tel: 800-553-4562 or 941-924-6565. Fax: 941-924-4892. For customer service, send email to Doni Richardson at admin@watsonwalker.com.

© 1999 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.

The following are trademarks of IBM Corporation: MVS/ESA, PR/SM, ES/9000, ESCON, CICS, DB2, DFSMS. All other product names are trademarks of their respective owners.

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

Cheryl's List

Here's a summary of the last transmissions sent to subscribers of our free electronic **Cheryl'sList**. 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 <<http://www.watsonwalker.com/archives.html>>. See the last paragraph below for instructions on how to sign up for the list.

Cheryl's List #17 - 10 December, 1998

In this issue, I'll cover the following:

1. Some Excitement in Our Life
2. Cheryl Watson's TUNING Letter, 1998, No. 5 Summary (*previously printed in issue 1998, No. 5*)
3. CPU Chart Mailed

1. Some Excitement in Our Life

First of all, we'll be moving to different offices here in Sarasota sometime in January. As soon as we finalize our schedule, address, and new phone numbers, we'll let you know. Our email ids will remain the same, as will our 800-number (800-553-4562). *Editor's note: Our new address and phone numbers can be found on page 4 of this issue.*)

But an even more disturbing event has taken place. Around noon on Monday Tom and I were at home when an armed burglar forced his way into our home. After sticking me in the trunk of our Lexus, he took our cash and jewelry and tied Tom up. He then told Tom that because we'd seen his face, he was going to have to kill us! When he left the room to look for more goodies, Tom escaped, ran to the neighbors, and called the police. So the robber took off with my BMW. It was recovered a short distance away with only minor damage. Our damage is a little less minor. Tom has a torn rotator cuff that will take several months to mend. (He had to leap the fence and run through some brambles.) He's my hero! We are dealing with the stress pretty well. We may move our personal residence, however, since the person is still at large. Because of all this, the last issue of 1998 will probably not be mailed until the latter part of January 1999. Our responses to email might be a bit slow also. We hope you understand.

3. CPU Chart Mailed

Cheryl Watson's CPU Chart was mailed on November 11, 1998. It contains the following changes:

- The IBM G5 series (9672-Rx6) announced May 7, 1998 has been updated with processor groups, SU/Sec, and our MIPS by workload. New turbo models that were provided in late October have been added.
- The HDS Pilot 98 series announced May 14, 1998 has also been updated and includes new processors announced on September 21, 1998. New Skyline models have been added.

- Amdahl's 40 new models of their latest Millennium 800 series (announced on October 12, 1988) have been added. A new Amdahl GS722 model was also added.
- Some processor groups were changed or added in the CPU Chart and are noted by a revision mark
- At the time that IBM re-calculated their LSPRs using an OS/390 software base, most of our customers were still running MVS. Therefore, for the older machines, we kept our estimated MIPS that were based on the older LSPRs. Now that most of our customers are on OS/390 (or will be very soon), we have changed those estimates to correspond to IBM's OS/390-based LSPRs. The result is a slight change (typically an increase) in estimated MIPS. If you are still running MVS, you may prefer to use our MAY 1998 Chart instead of this one.
- Rather than use IBM's LSPR values for the basis of our estimated MIPS by workload, we are now using a combination of LSPR values, the vendor's claims, and customer feedback. A comparison of workload MIPS should now give you a closer estimate to what you might expect to see for your workloads.
- We've added a new column, STIDP (store processor id), to show the actual processor model that appears in the machine and in your SMF type 70 data. The values for the older machines (over 8 years) have not been confirmed.
- The MIPS Basis column has been replaced with a column for HDS models to indicate the comparable IBM model (according to HDS).
- Many more models have been updated to include our estimated MIPS by workload.

All in all, there's information on 667 processors! If you don't subscribe to the TUNING Letter, you can still order one. The cost is just \$80 and includes any past issue of the TUNING Letter as well.

That's all for now. Stay tuned!

Cheryl Watson

=====

If you would like to obtain an email subscription to Cheryl's List, just go to our Web page <<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>. ■