MAINFRAME JOURNAL

For Users of IBM System/370 Architecture & Compatible Systems

November/December 1988

Joy Technologies Inc. Mining Machinery Division Rides Out Triple Conversion

THE PERFORMANCE ADVISOR

Now there's a doctor for your MVS system

Symptoms:

If treating an ailing MVS system isn't your area of expertise, let the **new** PC-based Performance Advisor Expert System serve as your MVS System's personal physician.

To begin with, Performance Advisor works with you to conduct periodic checkups, responding to symptoms observed in your mainframe. But it doesn't quit there . . .

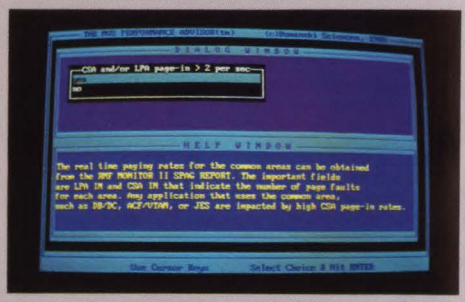

Diagnosis:

Performance Advisor proceeds to diagnose each problem and explain it to you in understandable terms so you know exactly what's wrong.

Soon you'll be able to diagnose these symptoms yourself. But don't worry, Performance Advisor is always on call for your convenience.

Treatment:

Once the diagnosis is complete, Performance Advisor prescribes a specific treatment. The result? This amazing new Expert System eliminates your MVS dilemmas by leading you through a simple three-step process that puts your system back into action!

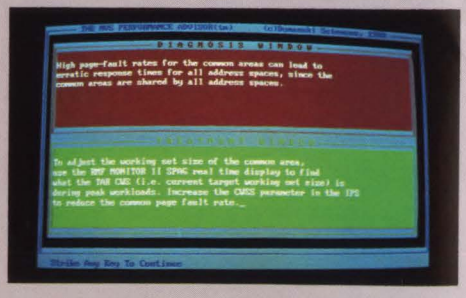

Also, customize your Knowledge BAses (KBA) by adding, modifying and deleting tuning rules of thumb according to your needs. You can even *share* knowledge by using our electronic Bulletin Board System.

Performance Advisor is also a valuable training tool that helps teach new people how to tune an MVS System.

Performance Advisor is just what the doctor ordered for maintaining your MVS System! Call Domanski Sciences today at 1-201-367-0257.

Priced at:

$895

DOMANSKI SCIENCES

16 Colonial Court, Howell, NJ 07731
(201) 367-0257

CIRCLE #55 on Reader Service Card ▲

Still screaming about why CICS failed?

Get the answer faster than you can print a system dump.

Eyewitness™ gives you **online diagnostics** instantly, so you can debug storage violations, CICS system abends, and operator cancels—fast. It even slashes dump time and downtime by up to 90%! There's no other product like it.

In just two months, Eyewitness helped over 75 data centers solve their CICS mysteries. Imagine…No more waiting for the dump to print. No more reams of paper. No more agonizing hours searching hex code for clues. And now, if you need level 2, you'll know it! Best of all, Eyewitness bears the signature of Landmark Systems Corporation, makers of The Monitor for CICS.™

Eyewitness is revolutionary! See for yourself. Call us today for a FREE, 30-DAY TRIAL at **1-800-227-8911** or 1-703-893-9046—and scream no more.

LANDMARK™

eyewitness™

SOLVE THE MYSTERY

Landmark Systems Corporation
8000 Towers Crescent Drive
Vienna, Virginia 22180-2700

CIRCLE #29 on Reader Service Card ▲

CONTENTS

November/ December 1988

Volume III Number 6

DEPARTMENTS

6 Publisher's Comments

8 Reader Forum

100 Product Review

106 Tech Advisor

108 Product Update

110 VSE Forum

114 Vendor Profile

COVER:

Joy Technologies Inc. manufactures underground mining equipment like the continuous miner that cuts coal pictured on the cover. The company recently underwent a triple conversion from:

- *DOS to MVS*
- *4381 to 3081*
- *DBOMP and IDMS.*

Photography is by Joy Technologies Inc.

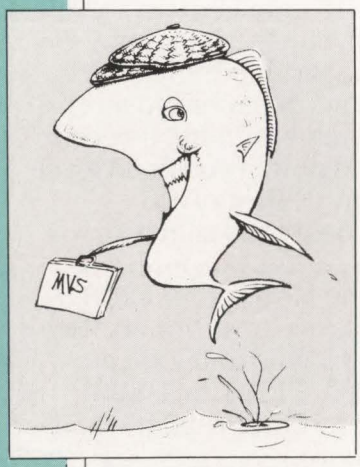

76 *Charley the Tuner*

FEATURES

10 VM/XA Comes of Age *By Romney White*

14 MVS Cross Memory Services: Share and Share Alike *By Bill Carico*

21 CICS Tuning Experiences *By Frank Bereznay*

29 ISPF Spells Productivity *By Jon E. Pearkins*

37 The Basics of ACF/VTAM Logmode Tables *By Lloyd A. Hagemo, Jr.*

40 Control DASD Growth Through Effective Storage Management *By Steve Adil*

46 VSE Under the Covers: A Look at the Internals of the VSE Operation System *By Eric Vaughan*

56 Design, Performance & Capacity Planning for DB2 Applications *By Joel Goldstein*

68 Strategies for Improved SQL/DS Performance *By Robert D. Smith and Lani Spund, PhD.*

72 Joy Technologies Rides Out Triple Conversion *By Lamont Wood*

76 MVS Performance Management Legends *By Stephen L. Samson*

88 Automation: A Managerial Discussion *By Greg Goticchia*

95 Strategic Alliance: Landmark Systems & Morino Associates Demonstrate Advantages *By Fred Newberry*

100 Product Review: VMOPERATOR Puts Muscle Into the Operator Console *By John Kador*

110 VSE Forum: Situations & Solutions *By Pete Clark*

114 Vendor Profile: SDI

MAINFRAME JOURNAL© (ISSN 0895-5751) is published bimonthly by Thomas Publications, Inc., 10935 Estate Lane, Suite 375, Dallas, TX 75238, (214) 343-3717. Second class postage paid at Dallas, TX. POSTMASTER: Send address changes to: MAINFRAME JOURNAL, P.O. Box 551628, Dallas, TX 75355-1628.

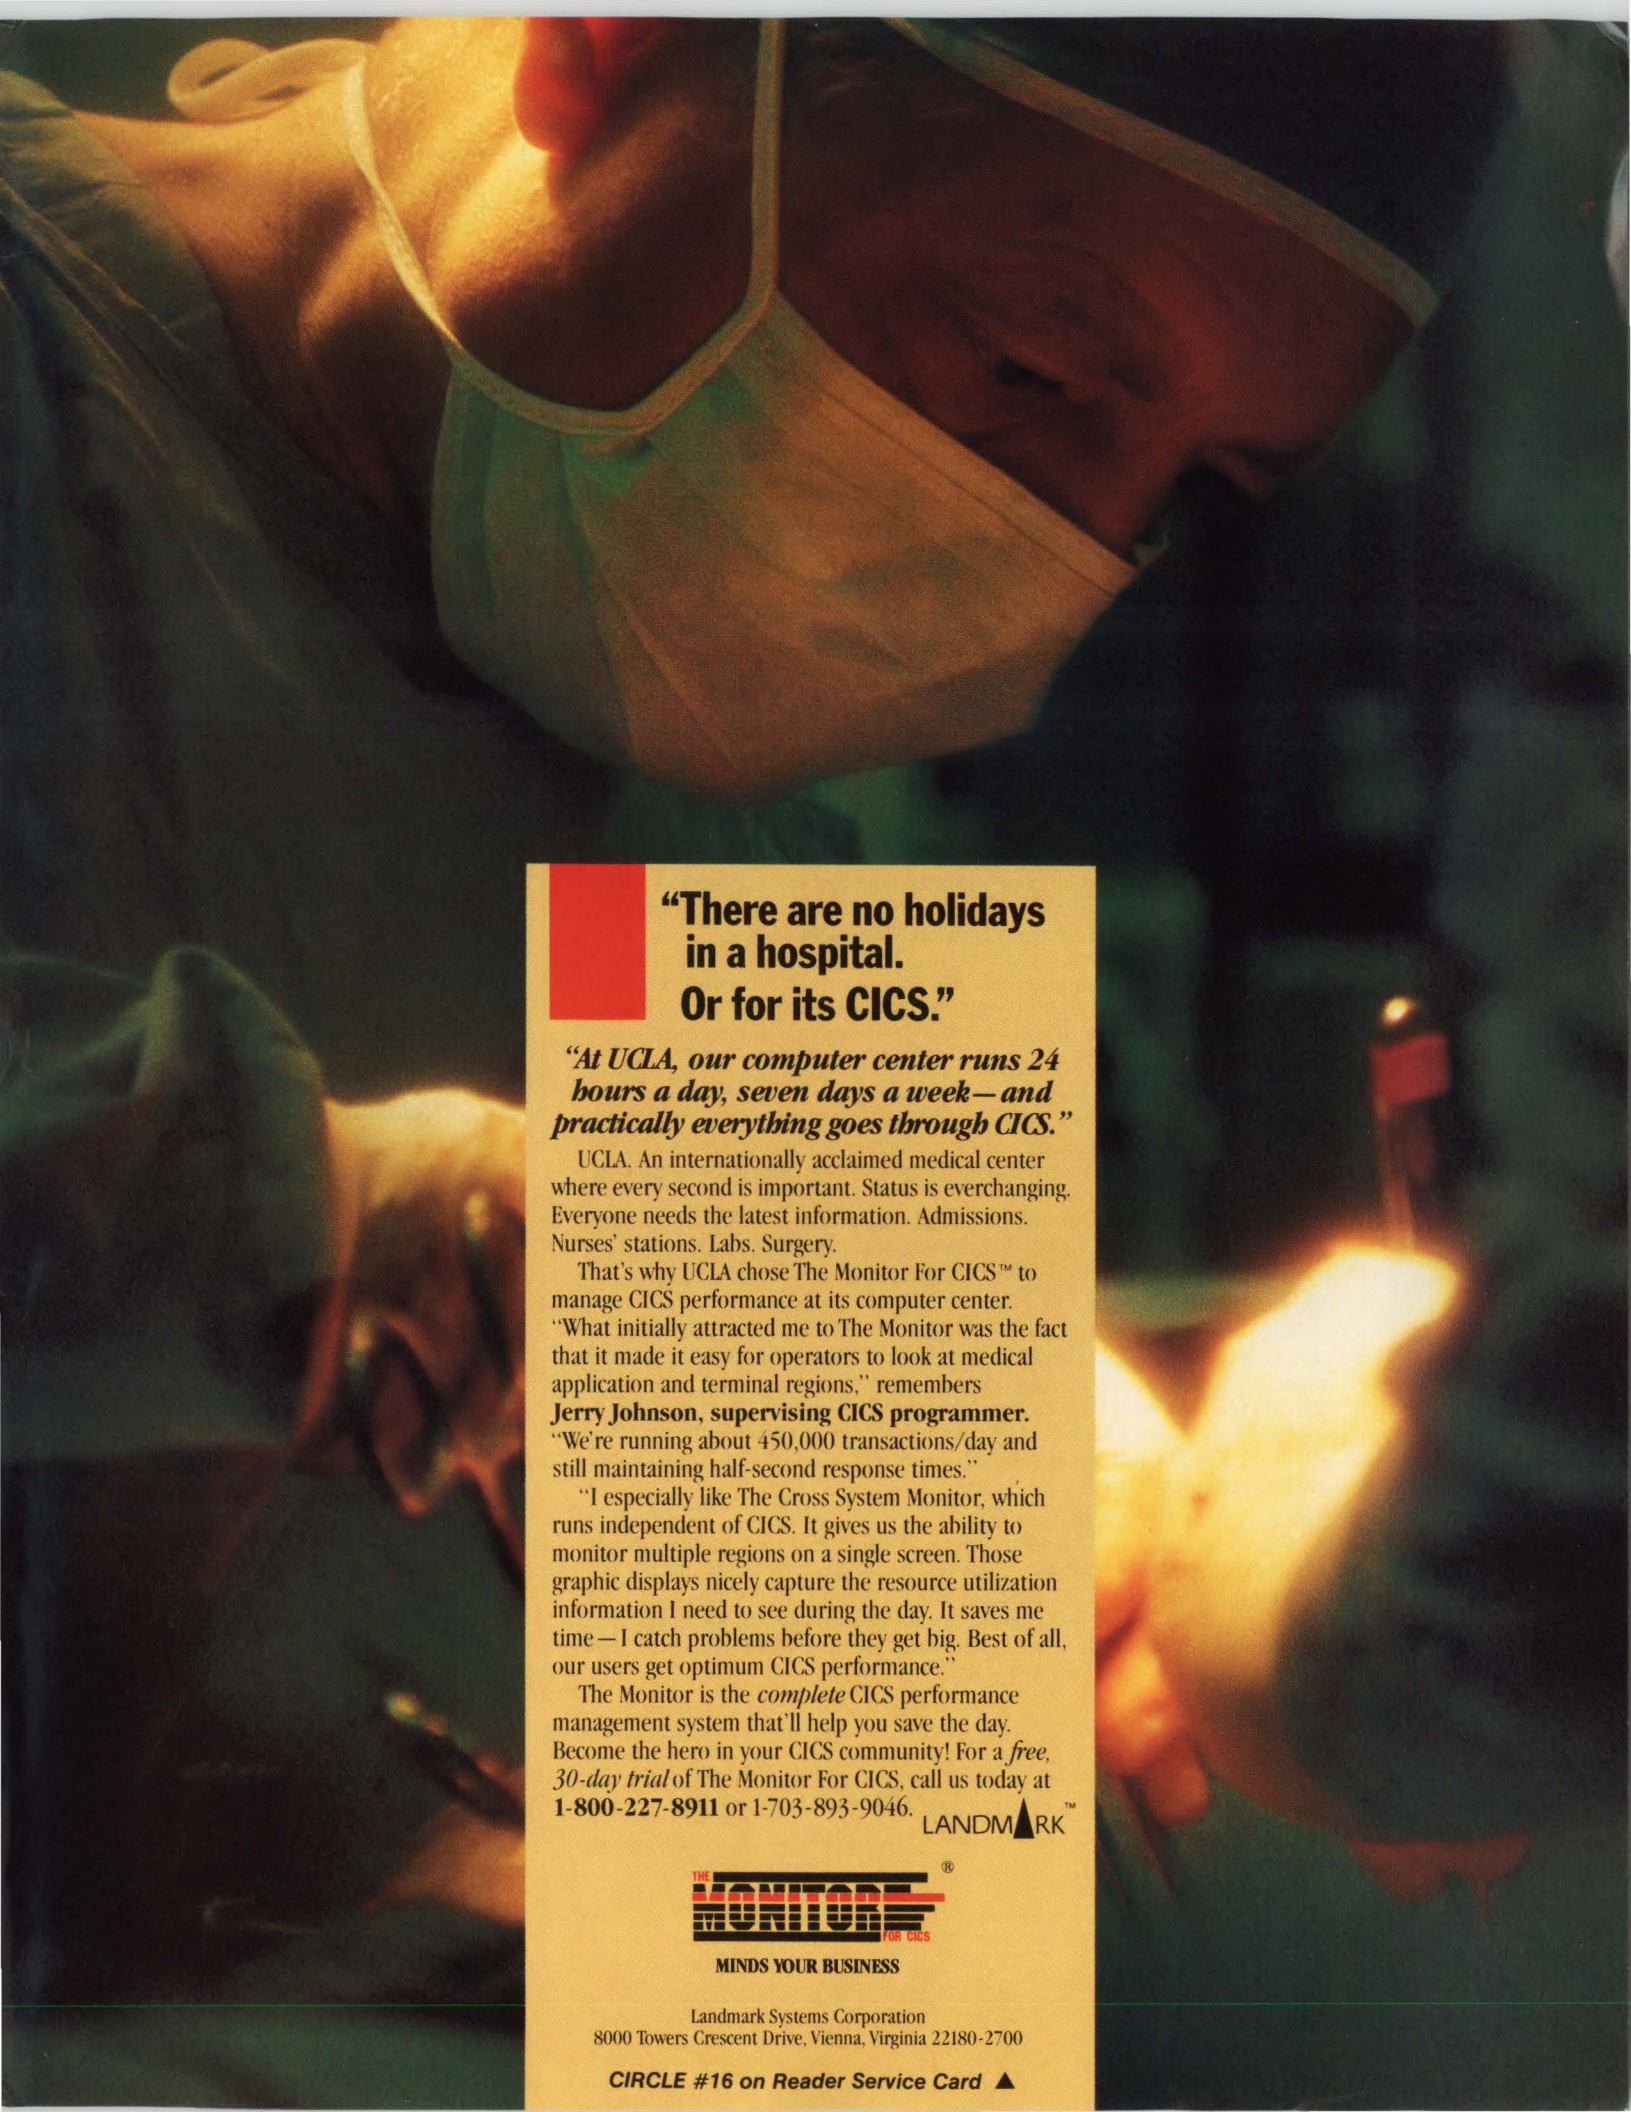

"There are no holidays in a hospital. Or for its CICS."

"At UCLA, our computer center runs 24 hours a day, seven days a week—and practically everything goes through CICS."

UCLA. An internationally acclaimed medical center where every second is important. Status is everchanging. Everyone needs the latest information. Admissions. Nurses' stations. Labs. Surgery.

That's why UCLA chose The Monitor For CICS™ to manage CICS performance at its computer center. "What initially attracted me to The Monitor was the fact that it made it easy for operators to look at medical application and terminal regions," remembers **Jerry Johnson, supervising CICS programmer.** "We're running about 450,000 transactions/day and still maintaining half-second response times."

"I especially like The Cross System Monitor, which runs independent of CICS. It gives us the ability to monitor multiple regions on a single screen. Those graphic displays nicely capture the resource utilization information I need to see during the day. It saves me time—I catch problems before they get big. Best of all, our users get optimum CICS performance."

The Monitor is the *complete* CICS performance management system that'll help you save the day. Become the hero in your CICS community! For a *free, 30-day trial* of The Monitor For CICS, call us today at **1-800-227-8911** or 1-703-893-9046.

LANDMARK™

THE MONITOR® FOR CICS

MINDS YOUR BUSINESS

Landmark Systems Corporation
8000 Towers Crescent Drive, Vienna, Virginia 22180-2700

CIRCLE #16 on Reader Service Card ▲

PUBLISHER'S COMMENTS

Robert H. Thomas

Recently, I attended the ADAPSO conference held in Dallas. ADAPSO is the computer software and services industry trade association. In talking with leaders in the software industry, I heard the term, "Strategic Alliance," used frequently. In fact, one whole session was devoted to this topic.

Strategic Alliances

The third-party software industry serving users of IBM mainframe systems has done an excellent job of providing a wide range of products designed to improve productivity and efficiency.

Most DP organizations now have a number of systems and application software packages from a variety of vendors. This disparity of packages and vendors has created a "Tower of Babel" effect. Typically, the software from the various vendors does not "communicate," leaving users with a fragmented collection of non-integrated packages or having to patch in homegrown routines to transfer data from package to package.

Fortunately this scenario is starting to change, much to the benefit of the users. Strategic Alliances are being formed whereby two (or more) vendors undertake a joint effort to create the necessary interfaces or bridges enabling their respective software products to operate harmoniously.

One example of just such a Strategic Alliance is chronicled starting on page 95. This article, featuring Landmark Systems and Morino Associates, provides insight into how and why the Alliance was formed and more importantly, how their customers reacted to this unique arrangement.

Triple Conversion Experiences

One morning several weeks ago, I received a call from a fellow in Pittsburgh, PA named Howard Horton. Howard simply called to say that his company (Joy Technologies Inc. Mining Machinery Division) had just completed a major conversion effort and thought the experiences might be interesting and helpful to others.

Read Lamont Wood's profile of Joy's *triple* conversion starting on page 72. Horton was exactly right when he suggested that the company's story might be helpful to others. I am grateful that he took the time to call.

Share Your Expertise

Much of the editorial content published in *MAINFRAME JOURNAL* is the direct result of readers taking the time to call or write offering to contribute an article on a particular topic of interest to our unique readership. This is great! We encourage anyone interested in sharing his or her expertise, a unique solution or a technical tip to either give me a call, use the enclosed Feedback Card or just submit an article outright. Contributors of published articles are compensated.

MVS — Not Babe Ruth

When you think of a "legend," your first thoughts may be of someone like Babe Ruth, Davy Crockett, Elvis Presley or some other famous person. Stephen Samson takes a far different look at legends in his article, "MVS Performance Management Legends," starting on page 76. In it, Samson debunks many of the entrenched beliefs about MVS that have reached legend status.

Robert H. Thomas

MAINFRAME JOURNAL
For Users of IBM System/370 Architecture & Compatible Systems

Publisher and Editor-in-Chief
Robert H. Thomas

Associate Publisher
Martha Thomas

Editor
Carol M. Hoag

Copy Editors
Judy Beller
Pat Warner

Art Director
Valerie Hall

Illustrator
Jeff Hill

Advertising Manager
Denise Thomas

Marketing Services
Sally Webb

Circulation Manager
Janice Porter

Assistant Circulation Manager
Nancy Crawford

Administrative Manager
Marian Davenport

Typography Coordinator
Dub Pope
Jaggars Chiles Stovall

Printing Coordinator
Tami Gustafson
American Signature Graphics

BPA Membership Applied For April, 1988.

SUBSCRIPTION RATES & INQUIRIES: Subscriptions are free within the USA and Canada. One-year foreign subscriptions are $96.

All subscriptions, remittances, requests and changes of address should be sent to MAINFRAME JOURNAL at 10935 Estate Lane, Suite 375, Dallas, TX 75238, (214) 343-3717.

MAINFRAME JOURNAL© is copyrighted 1988. All rights are reserved. Reproductions in whole or part prohibited except by permission in writing.

Syllogy announces an online sort that will bring down the house instead of the system.

CICSORT™

CICSORT FROM SYLLOGY. ONLINE SORTING FOR THE FIRST TIME EVER.

For years no one would even think of doing a sort online. Because calling the batch sort would cause CICS to crash. That's about to change. Introducing CICSORT™. A remarkable new technology that now makes it possible and practical to sort online. **FAST.** CICSORT lets you get critical reports faster than ever. No more waiting for batch reports. Transfer them online. Create new reports in CICS. Or upgrade unsorted reports. **EASY.** CICSORT is called by the standard COBOL Sort verb. Programmers can put it to work immediately. And it's fully compatible with the CICS preprocessor, the OS/VS and VS COBOL II compilers and all versions of CICS. **EFFICIENT.** CICSORT is designed to operate at peak efficiency under CICS. Without affecting the performance of other jobs. **PHONE.** Find out what CICSORT can do for you. Call to receive our free booklet about online sorting and ask about our 30-day free trial. **CALL 1-201-343-8900.**

syllogy
FOR INNOVATIVE SOFTWARE

SYLLOGY CORPORATION ONE UNIVERSITY PLAZA HACKENSACK, NJ 07601 (201) 343-8900

CIRCLE #85 on Reader Service Card ▲

CICS VTAM ACB Reopen Program

Here at Travelers we have a unique program (CICS VTAM ACB Reopen) that has saved us from restarting the nearly 50 production CICS regions after VTAM drop and restarts for more than a year now. I hope it helps other MAINFRAME JOURNAL readers as well.

CICS VTAM ACB Reopen Program

The Drop and Restart of VTAM will close the VTAM ACBs of CICS regions running on the same host. Customers will be unable to logon until it is reopened by some external event such as the console operator issuing the CEMT SET VTAM OPEN command. The worst thing that can happen is that this condition is diagnosed as a CICS problem and CICS regions on that host are also dropped and restarted. This can be a very time-consuming process. If automation tools are already in place in your data center, perhaps this process can be filled with the VTAM start procedures. However if this is not possible, a CICS transaction that reopens the VTAM ACB may be the best way to solve this problem. In this way, you can drop and restart VTAM and automatically have the CICS regions reopen their ACB themselves.

I have included a program that will accomplish this. It is a command-level, Assembler program and it is started in the PLTPI table. It checks to ensure that PLTPI processing is complete (you don't want to attempt any terminal control processing until PLTPI processing is done). The program links to DFHEMTD, the program invoked by the CEMT transaction, and inputs the CEMT SET VTAM OPEN command. After the link is complete, it issues an interval control start on itself so that it is executed every 60 seconds. With this program running every minute, each CICS region will have a 30-second mean time to recover after a VTAM drop and restart.

The CICS Gen requirements are simple and outlined as follows:

- The PPT must have a program entry, in this case MZ766ACB
- The PCT must have a transaction entry, in this case CACB
- The PLT specified by the PLTPI parameter in the SIT must have an entry, in this case MZ766ACB.

If you wish to stop the execution of this program, just disable the program via CEMT. This will cause the next start to fail and will stop the repetitive start process.

The performance of this program is very good. The total end-to-end transaction response time is under 0.4 seconds and it consumes under 50 milliseconds of CPU time. These measurements were taken on a 3084Q running in partitioned mode. This program was tested with CICS Releases 1.7 and 2.1.

Thomas M. Samulenas
The Travelers Companies
Hartford, CT

Editors Note: *Tom was kind enough to provide the code for his CICS VTAM ACB Reopen program. We will be happy to send it to any reader requesting it. To receive your copy, just CIRCLE #390 on the Reader Service Card.*

Obsolete Mystery Phases Problem Solved

In the July/August, 1988 issue of MAINFRAME JOURNAL, Warren Jackson of Washoe County School District described the problem of how to tag phases with the "last executed date" on his VSE/SP 2.1.5 guest.

We operate a VSE/SP 2.1.3 guest and we've developed a method of achieving this by using the history data of a job accounting package (in our case, CA-JARS). We dump the POWER account file daily (for subsequent merging and reporting) and also create a daily history file from which we pull selected up-to-date information and update a VSAM file. In turn, one of our reporting functions from this file is the last executed date and time of a phase. This report is cross-referenced with IBM's Library Directory Listing (giving first stored/last compiled dates) to make the phase entity complete.

Ottavio Constantini
Ayerst Laboratories
St. Laurent, Quebec

CIO Qualifications

I couldn't help but respond to R. Douglas Swords' article, "Was Your CIO An Accountant?" (September/October, 1988). Swords seems to be missing the boat when he says that some poor MIS executive is misunderstood when he reports problems of a "purely technical nature" to his chief executives and is "through no fault of his own" labeled "too technical." It is absolutely 100% his fault!

This guy should have the breadth to understand that his chief executives aren't going to understand or be interested in his behind the scenes "bits and bytes" problems. What they're going to be concerned about are the business consequences of those problems and if he can't relate that information he is wasting their time and his!

It takes more than an understanding of the technical aspects of MIS to be effective as a CIO. One needs a broad understanding of business and people to have any impact at that level.

Gary L. Wing
NBI, Inc.
Boulder, CO

CASE for Documentation

The article, "Documentation — What Documentation?" (September/October, 1988), that deals with the well-known problem of unavailable documentation for developed systems, misses a major point.

The author fails to mention the fact that the current market trend in CASE is to provide systems development tools that produce, often as an automatic by-product, documentation that represents the system being developed. The CASE method for producing documentation not only saves time and money, but also produces an error-free final product.

Sherry Marcy
Meta Systems
Ann Arbor, MI

HELP WANTED

Sunkist is currently running OSVS1 Release 7D, under VM, on a 4381 P2. All batch processing is handled by OSVS1.

Recently 3380s were installed with the plan to convert both VM and OSVS1 to 3380s.

The VM conversion went smoothly, but there are many roadblocks to running OSVS1 with Basic Programming Extensions.

I would love to hear from readers of MAINFRAME JOURNAL who could provide some guidance on the following:

- OSVS1 users of 3380s
- OSVS1 user groups
- Running a heavy batch production workload under VM ONLY.

Any assistance would be greatly appreciated.

Raymond Shawstad
Sunkist Growers, Inc.
P.O. Box 7888
Van Nuys, CA 91409
(818) 986-4800

JOBTRAC™

PRODUCTION CONTROL AUTOMATION

Today's Answer To Tonight's Problems.

Unparalleled ISPF Dialogue Productivity • On-Line Run Documentation
Forecasting & Simulation On-Line • Installation in Hours, Not Days • Operator Command Scheduling
No Dedicated Hardware or JCL Changes • Distributed Processing to up to 256 Separate Locations
Complete Dependency and Event Triggering • Sysout Capture and Archival Included
Automatic Operator Facility Provides Programmable Message Reply
Automatic Date Card and Override Generation • Free 30 Day Trial

Bennett Software, Inc. **1-800-JOBTRAC** P.O. Box 96694 • Houston, Texas 77213

© 1987, Bennett Software, Inc.

CIRCLE #136 on Reader Service Card ▲

VM/XA Comes of Age

It may well be the only form of VM IBM will support and enhance in a few years.

By Romney White

VM/XA, a version of VM that exploits the IBM System/370 Extended Architecture, has been in existence for quite a while. Until recently its use was limited to a small subset of the VM user community. This article will explain how this situation developed and what caused it to change.

VM/XA first saw the light of day as the VM/XA Migration Aid and was delivered by IBM to customers in 1983. The migration aid's primary target market was composed of large installations that wanted to migrate from MVS/370 to MVS/XA without having to acquire a second computer system. Using traditional VM techniques, this initial offering allowed several virtual machines to run together on a single system that was operating in Extended Architecture mode. Special emphasis was placed on running one virtual machine with especially good performance so that the production MVS system ran well.

The migration aid was superseded in 1985 by the VM/XA System Facility. This version was also aimed primarily at the large installation with interest in running MVS guest operating systems. VM/XA SF provided support for 3090 processors including a new hardware capability called the Interpretive Execution Facility. An associated new instruction, Start Interpretive Execution (SIE), was introduced to provide an invoking mechanism.

SIE took many of the simulation functions that VM usually performed and moved them under hardware (that is, microcode) control. Virtual machines could execute without frequent intervention from the VM Control Program (CP). Some people believed that SIE was a hardware implementation of CP. In fact, it has more in common with the Virtual Machine Assist microcode, available on many System/370 processors.

Recent Developments

In April, 1988 IBM introduced the VM/XA System Product (SP) and made it available to selected customers under a controlled release program. Unlike earlier VM/XA offerings, SP included a version of the Conversational Monitor System (CMS, VM's timesharing component) that exploited XA facilities. With this capability, a fundamental limitation inherent in earlier XA implementations of VM was removed. As a result, many more installations could consider moving to VM/XA. This version also included greater compatibility with its System/370 antecedents and provided basic VM facilities such as inter-machine communication and performance data collection.

Release 2 of VM/XA SP is scheduled to be delivered in late 1988. This version

Now, for your VM environment:

A complete framework for producing custom-tailored standards and procedures!

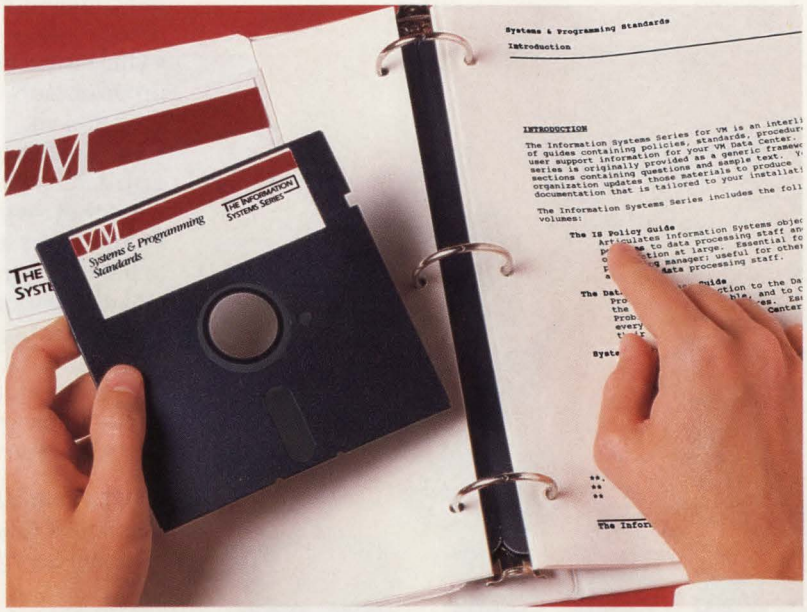

Here's what you've been waiting for! Now, with the ***Information Systems Series***™ documentation framework, you get everything you need to quickly develop custom policies, standards, procedures, and user support documentation for your VM installation.

That's right. Because you begin with a field-tested framework of proven documentation, you'll be able to cut down on time, effort, and the size of your documentation budget!

In addition, the ***Information Systems Series*** can help you implement IBM's **Systems Application Architecture (SAA),** as well as significantly expedite the process of IS **audit** approvals.

Here's what you'll receive.

The ***Information Systems Series*** is an interlinked set of four manuals. The accompanying diskettes, containing the very same material, allow you to personalize the documentation using your own word processing software on an IBM PC or compatible. It's like having a technical standards and procedures expert sitting right at your side, helping you customize the manuals so that they're always current and exactly right for your organization! The ***Information Systems Series*** will help you produce:

- **The IS Policy Guide.** It lets you clearly communicate Information Systems policies to staff and the organization, and helps maintain consistency in departmental activities.
- **The Data Center User Guide.** It contains information and procedures needed by everyone (new or experienced, technical or non-technical) to make effective use of the VM installation.
- **Systems & Programming Standards.** It defines installation-wide standards for the development and maintenance of production systems. The result? Easier-to-use systems and greater programmer productivity.
- **The Operations Guide.** It contains clearly-specified procedures for the operation of hardware, software, and the network, plus problem management and emergency procedures.

Call for more information.

It's easy to find out more. Just give us a call, at **800-777-0970** and we'll send you detailed information. There's no obligation of any kind.

The Computer Resources Group, Inc.

303 Sacramento Street, San Francisco, CA 94111

(Note: The Information Systems Series will soon be available for MVS!)

© Copyright 1988 CRG

CIRCLE #69 on Reader Service Card ▲

Why our system programming classes are full of people who use our competitors' CPUs.

It's true. More than two-thirds of our students work in organizations that use our competitors' machines. They come to us for five reasons:

- Our instructors are veteran system programmers who've practiced what they preach.
- They're also professionally trained instructors who know how to preach what they've practiced.
- Our classes are small, so it's easy to have the sort of give and take that makes learning fun.
- Many are taught in labs with on-line systems, so students can practice what their instructors preach, hands-on.
- They know you have to study with the best when you aim to be a grand master.

Choose from 50+ courses.

This year's Amdahl Education and Professional Services curriculum covers:

- MVS
- MVS/XA
- SMP/E
- VSAM
- JCL
- VM/CMS
- CP
- VM/HPO
- VM
- IMS/VS
- ACF/NCP
- ACF/VTAM
- SNA
- JES2
- ASSEMBLER

And you can take these courses in these cities:

- Atlanta
- Chicago
- Houston
- New York
- Columbia (MD)
- Los Angeles
- Orange (CA)
- Santa Clara (CA)

For a catalog that details our curriculum, call:
1-800-233-9521, ext. 60

amdahl®
The SMART Choice

CIRCLE #60 on Reader Service Card ▲

addresses another major requirement of large VM systems that previous XA systems did not — SNA support. With the ability to connect a VM/XA system to an existing SNA network, installations with large numbers of interactive terminals can start to make their way towards the XA environment.

XA Benefits

System/370 Extended Architecture offers a variety of benefits:

- Support for 31-bit addressing allows up to two gigabytes of main storage to be installed
- Expanded storage (also called paging storage) can be installed
- The architecture allows up to 64 processors — currently systems with from one to six processors are available
- Up to 256 channels are supported connecting up to 65,536 individual devices; the channel subsystem performs many of the functions necessary to start, manage and report input/output activity
- Bimodal addressing permits programs running in 24-bit addressing mode to coexist with those running in 31-bit mode
- Hardware trace facilities are provided.

VM/XA has exploited most of these architectural improvements in fairly straightforward ways. For example, it supports up to two gigabytes of main storage and up to six processors. It uses the hardware trace facilities to record significant Control Program events. However the more important consideration for many installations is how these architectural changes are used to expand the virtual machine environment. Some of the changes and improvements are as follows:

- Virtual storage sizes can range up to 999 MB
- Virtual machines can be defined with multiple virtual processors
- Both /370- and XA-mode virtual machines can be defined
- The new version of CMS can run in both /370 and XA mode earning it the name bimodal CMS; in XA-mode it allows both /370-mode and XA-mode programs to be run.

Implications

What all of these changes mean for VM installations is that it is time to start taking VM/XA seriously. Of course, if you operate a 4341 or a 9370, your hardware is not capable of running an XA system. The 4381 and 30XX systems can run VM/XA, however. Generally its capabilities will not benefit users of smaller 4381s, since many of the traditional assists for VM are not provided when XA mode is active. However installations with larger systems may want to consider installing VM/XA. There are potential benefits for them.

One benefit is the number of simultaneously logged-on users that can be supported will be greatly increased. The limit on the number of users supported by VM/HPO is about 1,200 due primarily to storage constraints below the 16MB main storage line. VM/XA eliminates this artificial limitation and may facilitate combining several separate VM systems into one.

Second, the XA input/output subsystem can deliver greater throughput with a given I/O configuration than VM/SP or VM/HPO. Also, improved expandability of the XA input/output configuration allows more resources to be applied to solving this traditional VM bottleneck. Performance measurement facilities provided by the channel subsystem permit the causes of I/O performance problems to be identified with greater accuracy. As a result, tuning and hardware upgrades can be more effective.

Third, CMS applications that require access to additional main storage can now be implemented. The limitation of 16MB of virtual machine storage has been removed.

Fourth, guest virtual machine performance continues to be provided with dedicated processor support, dedicated expanded storage facilities and support for several virtual machines with dedicated main storage.

Fifth, improvements to performance monitoring facilities permit a better understanding of what resources and services are constraining system and virtual machine performance. Information about transaction service times, device response, inter-machine communication and other aspects of system performance makes tuning easier and upgrade decisions less error-prone.

As with any new system, there are reasons to approach VM/XA with some degree of caution.

One reason is compatibility between VM/XA and VM/SP is not complete. Applications that rely on command response formats, message numbers or some CP or CMS facilities may not function correctly.

Another is some devices are not supported including line-mode terminal devices and 3370 DASD.

Third, most CMS applications do not exploit Extended Architecture and many do not function in an XA-mode virtual machine. Some products do not operate under bimodal CMS in either mode.

Fourth, migration between VM/SP and VM/XA is not convenient.

Fifth, a new, more complex, less useable installation and service mechanism has hampered system maintenance efforts. The quantity of Object-Code-Only (OCO) materials in CP and CMS has increased impeding problem diagnosis and repair and limited the system's extensibility.

Last, until there is a critical mass of customers using VM/XA in true production environments with CMS, the anticipated performance gains cannot be confirmed.

Many of the problems with software applications will be addressed in the relatively near future as vendors become more aware of VM/XA and users desiring to run it. The user community will grow to a respectable size as increasing numbers of installations recognize the value of the XA environment for meeting their interactive computing needs in addition to their guest system requirements.

Conclusion

VM/XA is about to arrive as a reasonably full-fledged VM system. Within a few years, it may well be the only form of VM being supported and enhanced by IBM. Every large VM installation should be considering how and when to move to VM/XA.

ABOUT THE AUTHOR

Romney White has been working with VM for 16 years. He recently co-founded Velocity Software, Inc., a company that provides XA products and services for VM users. Velocity Software, Inc., 60 Alban St., Boston, MA 02124, (617) 825-3599.

MVS Cross Memory Services

Share and Share Alike

By Bill Carico

Cross Memory Services (XMS) is an extension to the hardware architecture that MVS now supports allowing enhanced communication between address spaces.

Since it is geared for a multi-address space design, its use will typically be confined to MVS components and MVS subsystems. Two previous articles in this series of three discussed XMS and explained how XMS protection and authorization mechanisms work. This article will examine the macros used to establish the Program Call (PC) environment and the linkage tables used during the execution of the PC instruction.

As you might remember, XMS is organized into three broad categories: program sharing, data movement and data sharing. In program sharing, a program in one address space can issue a program call instruction to effectively "branch" to another program located in a different address space. This is accomplished using a PC instruction. Data can be moved directly between two private areas in data movement. In data sharing, a program running in a system common area can reference data located in two different address spaces by alternating between them.

Program sharing is by far the most complicated of the three features. In today's technical journals, articles that deal primarily with how a computer instruction works are seldom seen. However since the PC instruction is more like a subroutine than an instruction, it is worthy of special consideration.

At a high level, the PC instruction actually performs the same type of functions as the MVS dispatcher before passing control to a program. The difference is that the PC instruction does all of its magic within the system's microcode. By reading further you will discover how PC routines get set up initially, how late binding is accomplished, how system controls and status indicators are changed by a PC instruction, how the PC instruction has been enhanced for MVS/ESA and how PC routines are removed and cleaned up.

Setting Up a PC Routine

The program sharing capability was provided in order to allow MVS components and subsystems to place their supervisory and/or service programs in the private areas of their respective address spaces while still providing convenient access by programs running in different address spaces. This design reduces the requirements for common storage (provides virtual storage constraint relief) and bypasses the overhead of making several passes through the dispatcher code in order to alternate program execution between two address spaces.

In order to set up a PC routine that can be called from a different address space, information about that routine must be passed to the PC/AUTH address space. PC/AUTH is like a park ranger who will not let anyone occupy a campsite until he is properly registered. PC/AUTH services are invoked using macro instructions that are restricted to authorized programs (supervisor state or PSW key 0-7). Naturally whenever PC/AUTH macros are executed, they generate a PC instruction which in turn invokes a PC routine residing within the PC/AUTH address space.

The registration process is necessary for a couple of reasons: for security purposes, since PC routines can run in supervisor state and because every routine must have a unique number in order to be invoked by the PC instruction.

It is also important to understand that this "registration process" must be con-

ducted every time XMS initializes, which occurs whenever MVS is IPL'ed. This means that the numbers that are assigned to PC routines are never permanent and may theoretically change or differ from one IPL to the next. Whether or not the numbers actually change depends on whether or not the start up sequence changes for those address spaces that set up PC routines. PC numbers are assigned on a first-come, first-serve basis.

Linkage and Entry Tables

The first part of the registration process is accomplished by reserving an entry in the Linkage Table (LT). That is accomplished by calling PC/AUTH using the LXRES macro (Linkage indeX REServe). PC/AUTH assigned a linkage index number to the caller that in effect holds a place in the LT.

The second step is to create an Entry Table (ET) that will contain detailed information about the routines that will be later accessed by other programs using PC instructions. The Entry Table CREate (ETCRE) macro is issued by the caller and PC/AUTH builds an ET in its own Local System Queue Area (LSQA) that contains the program descriptions. ETs are built in increments of four entries, each called an Entry Table Entry (ETE). ETs may contain a maximum of 256 entries. The Entry Index (EX) of a specific routine is determined by the system or subsystem that owns the entry table.

The third step is for the caller to connect its entry table to its reserved location in the linkage table. This is done using the Entry Table CONnect (ETCON) macro. Once this is accomplished, the PC number for a given program in the ET can be determined by combining the LX number and the ETE number.

If you understand Assembler language programming, then all I have to say is that the LX and EX values are OR'ed to derive the PC number. For example, assuming that program XMS129C is the twelfth entry in the ET, its ETE number is x'0B' (the first ETE number is x'00'). If the assigned Linkage Index (LX) was x'00A', then XMS129C's PC number will be x'00A0B'. Likewise, if after the next IPL the LX assigned is x'009', the XMS129C would then be called using PC number x'0090B'.

The AX macros are used to establish correct authorization levels for using the SSAR and PT instructions. Since the PC routine will be called from other address spaces, the Program Transfer (PT) instruction will be used to return control to the caller. Setting the AX to '1' using the ATSET macro will provide PC routines with the authorization to issue the PT instruction. AX=0 is the default authorization index for an address space.

Late Binding

Because the LX is assigned dynamically, the PC numbers for MVS components are not permanently associated with PC routines in the same manner as MVS' Supervisor Call routines (SVC) numbers.

Since PC numbers are not known before program execution, macro instructions and other interfaces will not be able to use PC numbers directly but rather indirectly. This means that the subsystem providing service using a PC routine must use an indirect mechanism to let its potential callers find out which PC number to use. Usually this means passing the PC

NOW, CICS AND VSAM UNDER VM WITHOUT A GUEST OPERATING SYSTEM!

With VMCICS™ and VMVSAM™ from Unicorn Systems Company, you can develop and run your production CICS and VSAM applications directly under VM. And you can do so without compromising compatibility or performance. No resource hungry guest operating system is required. No hardware restrictions—from the IBM® 9370 to 43xx to 30xx and compatibles. No headaches.

VMCICS/Development System allows you to develop, test and debug CICS applications directly under VM with exceptional productivity. You can obtain true VSE or MVS CICS compatibility with the advantages of developing under VM/CMS. Your developed applications can be moved back to VSE or MVS for production or remain under VM using VMCICS/Execution System.

VMCICS/Execution System is a multi-user production CICS environment for VM which provides outstanding performance. CICS/VSAM applications easily port from VSE or MVS. And, under VMCICS/ES, they operate with improved stability—no more region crashes. In effect, users get their own "virtual region."

VMVSAM is a multi-user, shared file system for VM which provides full VSAM compatibility. Programs written in COBOL, Assembler or REXX can now share VMVSAM files. And VMVSAM supports concurrent sharing of files between batch and online programs operating under VMCICS/ES—with full data integrity.

Together, VMCICS and VMVSAM make it possible to move your CICS/VSAM applications to VM, the most popular and fastest growing IBM mainframe operating system. So isn't it time for you to say "no more guests"? Be *our* guest by calling toll-free today at **800-222-6974** (from California, **800-232-CICS**).

unicorn

Unicorn Systems Company
3807 Wilshire Blvd.
Los Angeles, CA 90010 (213) 380-6974

IBM Business Partner
Authorized Application Specialist

VMCICS and VMVSAM are trademarks of Unicorn Systems Company.

CIRCLE #44 on Reader Service Card ▲

number(s) as part of a parameter list, part of a common data area or as a record in a common dataset.

This requirement for having an interface mechanism for passing PC numbers prevents programs from having any PC numbers binding to their services. Late binding means that the caller of a PC function is never dependent upon the actual PC number, which module will perform the function or the location of that module.

In addition, the dynamic assignment of PC numbers has at least two other advantages over a static assignment. First, it keeps the size of the LT as small as possible. The LT's initial size is 128 bytes, which gives room for 32 entries (LX values). If a larger LT is required, it will be dynamically expanded in increments of 32 additional entries up to a maximum of 1,024 entries. Second, dynamic assignment of PC numbers avoids having to reserve a set of numbers for use by the installation.

The System Function Table

MVS uses a System Function Table (SFT) to contain the PC numbers that are assigned to system services. Each PC function provided by MVS has a permanently assigned index into the SFT that is used to save the PC number once it is assigned. After the PC number has been stored, the PC routine is invoked by macro logic that indexes directly to the fixed offset into the SFT in order to extract the actual PC number (see Figure 1).

To simplify this explanation, just think of the SFT as being a bulletin board where the daily PC numbers get posted for each day's use. Since the macros imbedded within calling programs always derive the PC number by looking at the same spot on the bulletin board, the PC programs are always readily accessible once they have been initialized.

A Dispatch from Microcode

In order to invoke a PC routine, the caller will most likely issue a macro instruction that expands into a PC instruction. As mentioned previously, the macro expansion must obtain the PC number through whatever mechanism has been established in order to include it as the operand of the PC instruction. Once executed, the PC instruction is designed to traverse the LT (using the LX portion of the PC number) and the ET (using the EX portion of the PC number) in order to isolate a single ETE (see Figure 2). The ETE contains the necessary information to switch to (pass control to) the PC routine.

An ETE contains the:

- Authorization Key Mask (AKM)
- Execution Key Mask (EKM)
- Instruction address of the PC routine

FIGURE 1

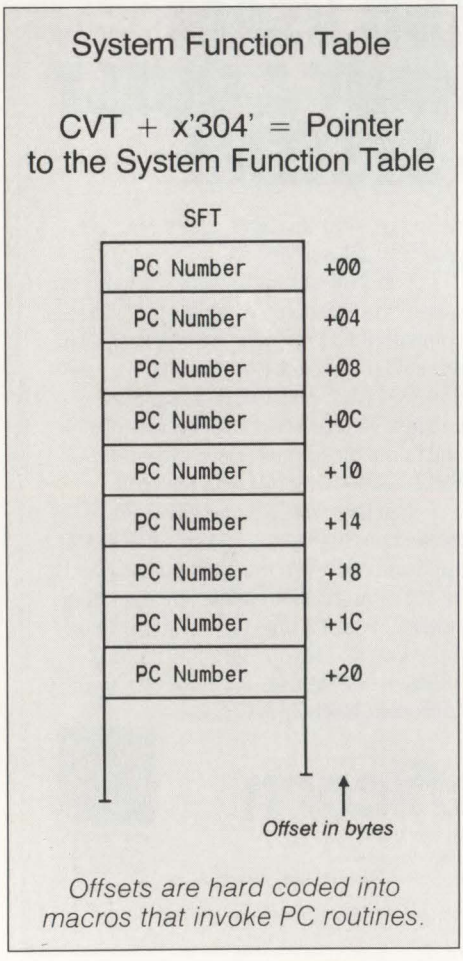

Offsets are hard coded into macros that invoke PC routines.

FIGURE 2

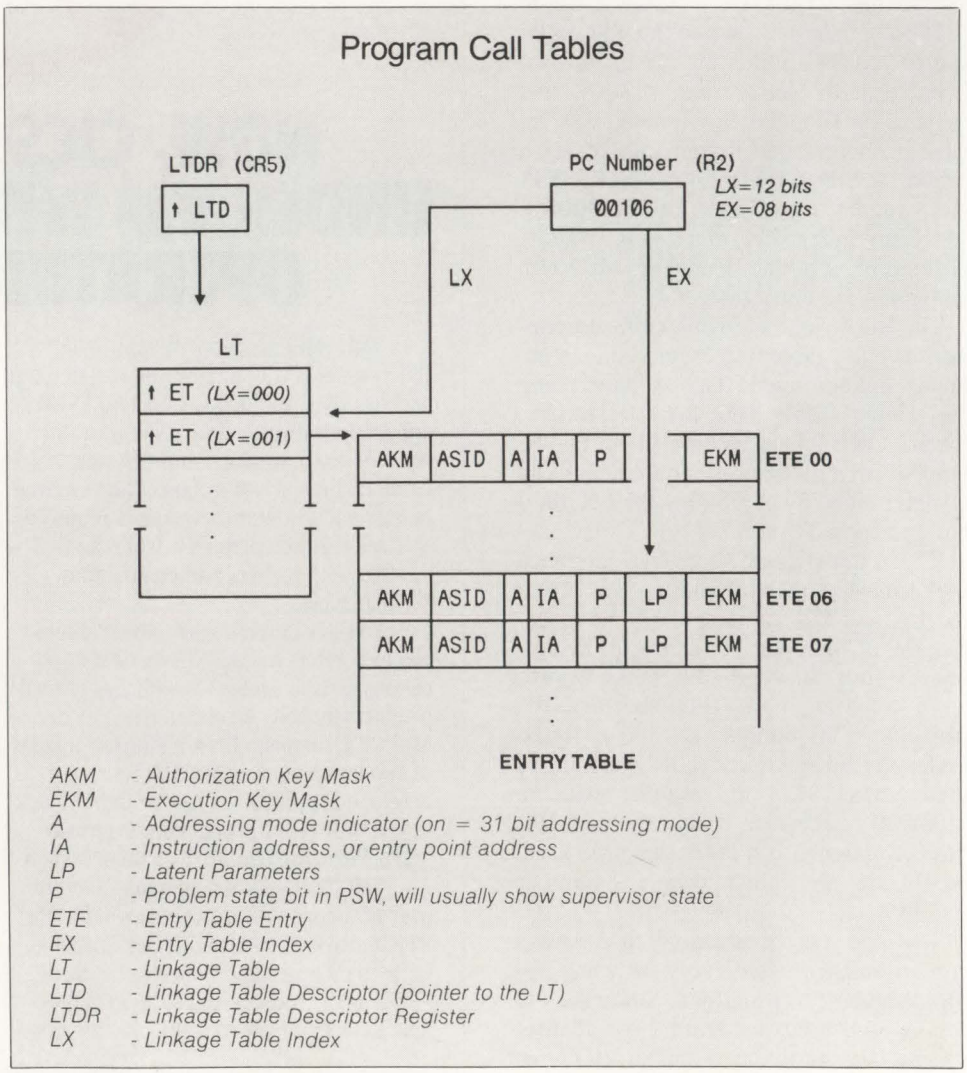

- PSW Addressing Mode (AMODE) indicator
- Bit 15 for a new PSW
- Address Space Identifier number (ASID) where the program will execute
- Pointer to parameters (called Latent Parameters — LP) that will be passed to the PC routine.

The hardware's microcode for the PC instruction uses this information to make changes to the system environment and then passes control to the program referenced in the ETE. The following discussion will elaborate on the details of how this works.

PSW Key Mask Manipulation

The AKM exists as a field in the ETE and is used to control access to the program that is pointed to by the ETE. When a problem state PC is issued, an authority check is performed to ensure that the program issuing the PC is in the correct state and key to call the desired function.

When an ETE is referenced by the PC instruction, the AKM field is AND'ed with the user's PKM (which is in Control Register 3). In the event no match is found, a program check interrupt occurs and the user is abended.

As the PC instruction continues execution through microcode, the PKM is updated (OR'ed) with the EKM to establish a new PKM for the PC routine. For example, a PKM of x'0080' when OR'ed with an EKM of x'8000' would result in a new PKM of x'8080'. Since each bit corresponds to a protect key, a value of x'8080' means that the PC routine can execute in either key 0 or key 8. (A detailed explanation of the PKM may be found in the previous article.)

Getting Addressability

The status quo for MVS/XA is to have addressability to two address spaces simultaneously using Control Registers (CR) 1 and 7. CR1 points to the segment table for the primary address space and CR7 is used for addressing a secondary address space. Under normal program execution, CR1 and CR7 point to the same segment table.

For a PC instruction, the microcode will update CR1 to point to the segment table for the new address space. Therefore, as a result of the PC being issued, the system establishes a new primary ASID. CR7, which is not changed, points to the home ASID.

Just Like an LPSW — Almost

The MVS dispatcher passes control to ready work using the LPSW instruction. LPSW replaces the entire 64 bits of the current PSW with the contents of a designated storage location. The PSW that is loaded usually changes the system state from disabled to enabled (for interrupts) and from supervisor state back to problem state. The instruction address of the PSW points to the program that will receive control and the Addressing Mode (AMODE) bit is set to indicate 24-or-31-bit mode. The PC instruction logic accomplishes a portion of the function provided by LPSW. The PC instruction may change the PSW (bit 15) to ensure that the PC routine receives control in supervisor state. It will update the instruction address to point to the PC routine and the AMODE bit is set appropriately. Note that the PSW interrupt masks are not affected by a PC instruction.

The VSAM Mechanic™

Now you can recover catalogs and VSAM datasets quickly even when other options fail. VSAM Mechanic will fix a broken VVDS in minutes without restoring a whole pack and backleveling users' costly data. Mechanic provides diagnostic functions to identify problems and commands to fix whatever is wrong. It can: ☐ automatically resynchronize a catalog with its volumes ☐ reconstruct catalogs & datasets that have been partially destroyed ☐ move catalogs to different volumes or to different device types ☐ change a catalog's dataset name ☐ rebuild free record chains ☐ delete CAXWAs ☐ delete "orphaned" components. VSAM Mechanic is a powerful tool for repairing ICF/VSAM catalogs and VSAM datasets.

For more information call
800-638-9254
MVS only.

People
SOFTWORKS, INC.

CIRCLE #179 on Reader Service Card ▲

Latent Parameters

The final phase of the instruction is to pass a pointer to the PC routine's LP. This pointer is part of an ETE that is stored when the ETE is initialized and points to a two-word parameter area. The PC instruction logic picks up the pointer and places it in register 4.

MVS/ESA Enhancements

As if this were not already enough for the PC instruction to do, PC has been enhanced as part of the new ESA/370 architecture. Previously, before a program issued a PC instruction, it would also have to save the environment (certain registers, PSW and PKM information using a PCLINK macro) and restore them when the PC routine returned control. The enhanced version of the PC instruction allows the system to automatically save the environment in a new hardware-controlled stack and to automatically restore the environment before returning control to the caller.

In addition, in order to save the overhead of establishing a recovery environment each time a PC routine is entered, the new PC instruction allows an associated recovery routine (similar to an ESTAE) to be set up as part of the PC instruction logic. The details of how these changes may have affected the ET format were not available at this time. In any case the PC instruction has been made even more powerful in ESA and many of the gyrations that a calling program had to go through when using PC have been eliminated.

Removing Access and Cleaning Up

The steps for removing PC routines (ETs) are essentially the opposite of the steps required for setting them up. Disconnecting the ET from the LT uses the ETDIS macro. The ET Destroy (ETDES) macro will clean up the ET and the LXFRE macro will free the LX so that it can be resued by another subsystem. Finally, AXSET and AXFRE will restore authorization tables and indexes to their previous state.

Formatted Dump Output

The example in Figure 3 is taken from an MVS/XA 2.2 SVC dump and represents some of the PC routines available in the system at the time the dump was taken. This figure only represents a portion of the actual entries formatted. As you can see, the dump formatting program has already combined the LT and *all* of the ETs into one large table rather than formatting each table separately. This format makes it much easier to access table information for diagnostic purposes.

FIGURE 3

Sample Table from an SVC Dump

PC NUM	AUTH KEY MASK	EXEC ASID	ENTRY ADDR	EXEC STATE	LATENT PARMS	EXEC KEY MASK
000000	FF00	0002	82201D18	S	7FFFE620	8000
000001	FF00	0002	822026B0	S	7FFFE628	8000
000002	FF00	0002	822033B0	S	7FFFE630	8000
000003	FF00	0002	82204190	S	7FFFE638	8000
000004	FF00	0002	82204930	S	7FFFE640	8000
000005	FF00	0002	82205E48	S	7FFFE648	8000
000006	FF00	0002	82206D00	S	7FFFE650	8000
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
000100	0000	0004	81131F9E	S	019F2A60	8000
000101	0000	0004	81132064	S	019F2A68	8000
000102	0000	0004	80008920	S	019F2A70	8000
000103	0000	0004	80008490	S	019F2A78	8000
000104	FFFF	0004	80009BD0	S	019F2A80	8000
000105	0000	0004	80008020	S	019F2A88	8000
000106	0000	0004	80009BD8	S	019F2A90	8000
000107	0000	0004	811329DE	S	019F2A98	8000
000108	0000	0004	8113310E	S	019F2AA0	8000
000109	0000	0004	80009660	S	019F2AA8	8000
00010A	0000	0004	80FFC1DC	S	019F2AB0	8000
00010B	FFFF	0000	8101A814	S	00000000	0000
000200	C000	0008	82200B50	S	019C6BC8	C000
000201	FFFF	0000	8101A814	S	00000000	0000
000202	FFFF	0000	8101A814	S	00000000	0000
000203	FFFF	0000	8101A814	S	00000000	0000
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.
000D00	FFFF	0042	00008C60	S	00FA2E68	FFFF
000D01	FFFF	0042	00008C60	S	00FA2E70	FFFF
000D02	FFFF	0042	00008C60	S	00FA2E78	FFFF
000D03	FFFF	0042	00008C60	S	00FA2E80	FFFF
000D04	FFFF	0042	00008C60	S	00FA2E88	FFFF
000D05	FFFF	0000	8101A814	S	00000000	0000
000D06	FFFF	0000	8101A814	S	00000000	0000
000D07	FFFF	0000	8101A814	S	00000000	0000
000E00	FFFF	004F	8001CDF8	S	0196A2C0	FFFF
000E01	FFFF	0000	8101A814	S	00000000	0000
000E02	FFFF	0000	8101A814	S	00000000	0000
000E03	FFFF	0000	8101A814	S	00000000	0000

Summary

This article has examined the details of how cross memory's program sharing feature works. The PC/AUTH address space provides control of the environment via a series of macros that can be issued by supervisor state, key 0-7 programs. The two main tables used to support program sharing are the LT and the ET. Once a provider of a PC function gets all of the tables in place, the PC number that corresponds to each routine listed in the ET can be passed to programs in other address spaces that will invoke the PC routines using a PC instruction.

This process of not having PC numbers hard coded in calling macros allows late binding to occur between service users and service providers. Once the binding is complete, the PC instruction is issued and it basically accomplishes the equivalent of a dispatch from within the hardware's microcode.

Parallels to the dispatcher are made due to the numerous similarities. The end result of a PC instruction is that a PC routine can receive control in a different address space using a different protect key mask, switch from problem to supervisor state and even change from 24-bit addressing mode to 31-bit mode or vice versa. Whether or not any of these status controls actually changes is determined by the contents of the ETE that exists for each PC service routine.

ABOUT THE AUTHOR

Bill Carico is president of ACTS Corporation specializing in providing systems training for the MVS environment.

Your Problem: Eliminate complex, time-consuming access to multiple applications and speed data exchange between them.

Your Solution: TUBES — the world's best session management tool.

- Quick to install, easy to administer
- User friendly, providing superb control and system security
- Available in any environment: VSE, VM or MVS with VTAM or BTAM
- Instant switching between 12 concurrent sessions with no loss of application context
- Transparent access to applications on any CPU

Call today for complete information or to start your no-obligation trial. Let **TUBES** deliver the "knockout punch" to your network performance problems!

MACRO 4... we've been providing powerful solutions to MIS system problems on an international scale since 1968.

MACRO 4... we deliver an unmatched level of technical expertise that has resulted in a full range of software solutions for VSE, VM and MVS operating system requirements.

ONE CALL TO 800-223-0414 PUTS THE POWER OF MACRO 4 IN YOUR CORNER!

Other MACRO 4 Performance Products:

VSAM/TUNE	The leading VSAM Performance Tool
LOGOUT	Console Manager for VSE and VM
CICSPRINT	Spooling Package for CICS, JES and Power
VPAC	VM Performance Analyzer and Control Facility
DUMPMASTER	Premier Dump Management and Analysis Package

macro 4
Systems Software
Brookside Plaza, P.O. Box 187, Mt. Freedom, NJ 07970 • (201) 895-4800 • (800) 223-0414

CIRCLE #39 on Reader Service Card ▲

Keep all your CICS regions out of the red.

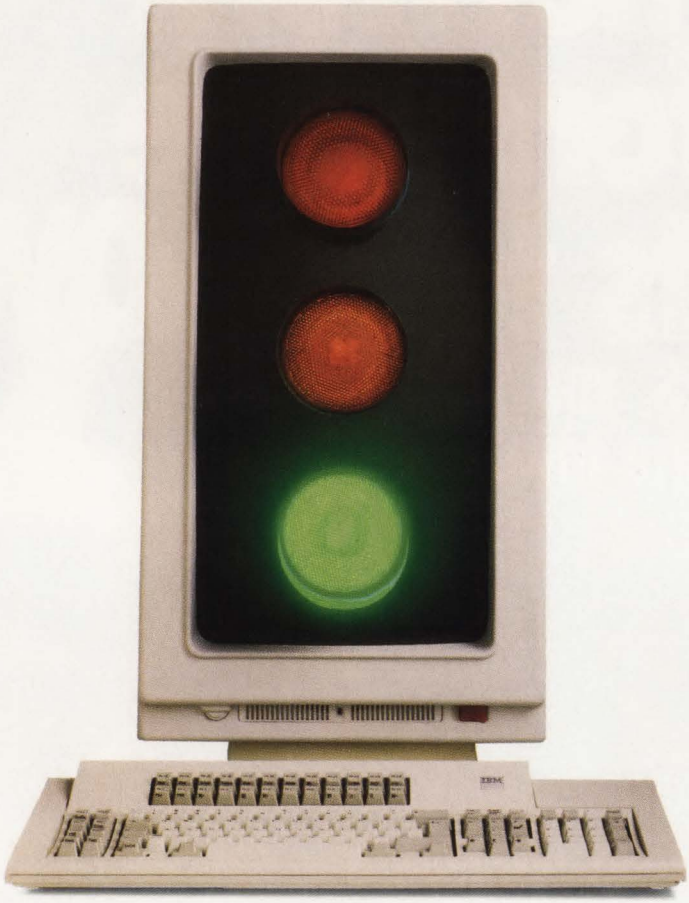

Managing your CICS regions can be as hectic as directing rush hour traffic. Sometimes just keeping track of a few critical production regions is a challenge.

But with Candle's OMEGAMON® for CICS, you can arrest an availability threat as it comes speeding around the corner. On a single screen, OMEGAMON's Status Monitor gives you a traffic 'copter's view of your entire CICS network — every region, on any CPU, in any geographic location. And the Status Monitor is free to all users of OMEGAMON for CICS.

The three-color display shows you at a glance where OMEGAMON is warning that a region's availability is in danger, where response time is stalled, or where DASD access is slower than acceptable.

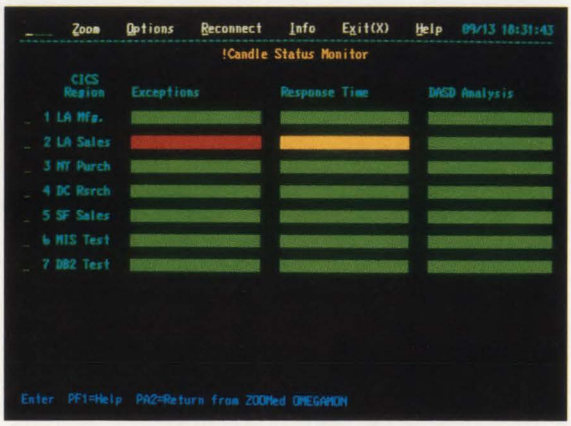

Monitor up to 34 regions on a single screen.

With instant access to the extensive problem-solving power of OMEGAMON, you can improve service in your critical regions by freeing up stalled traffic as soon as it appears. Before your users even notice a slowdown, you've pushed the transaction block out of the way and are well into clearing up the next bottleneck.

And whenever you need help. Candle's expert support team is always available. With round-the-clock customer service, technical education, and a commitment to stay current with IBM. So you won't ever get caught in CICS gridlock.

To find out how to keep all your CICS regions out of the red, call Terry Forbes today at **(800) 843-3970.**

!Candle®

CIRCLE #116 on Reader Service Card ▲

CICS Tuning Experiences

By Frank Bereznay

The Automobile Club of Southern California provides its 3.4 million members with a wide variety of motor club services and insurance products from more than 80 district offices located throughout Southern California. To support the delivery of these products and services, a network has been developed to connect more than 2,500 terminals located in the district offices to a large scale IBM system located at the Club's processing center in Costa Mesa. This IBM system is used to run a CICS workload to support the membership and insurance lines of business.

The processing complex at the beginning of the project consisted of two equally configured 3081-GX processors with 24MB real storage and 16 channels each. These processors share an I/O subsystem that consists of five 3880 storage controllers and 3380 series DASD. DASD devices were typically dedicated to storage pools based upon type of use but there was no assignment of channels based upon usage type or system. One of the processors, the on-line processor, was dedicated to the production CICS workload and the second processor was used as a development machine by the programming department and served as a ready backup to the on-line processor.

The CICS workload consisted of four regions. Two of the regions were dedicated to single applications with dedicated terminal populations, a system to dispatch tow trucks and a system used for data entry applications. The remaining two regions were using a Multiple Region Option (MRO) connection to support a collection of applications for district office operations. The MRO configuration had terminal definitions and workload in the primary region and a single significant application in the second region. This configuration was more a reaction to a Virtual Storage Constraint Relief (VSCR) problem than a planned MRO configuration.

Exploiting the MRO feature was the most significant change implemented and it provided a number of benefits.

Overall the CICS workload was delivering 835 transactions per minute with an average host response time of .93 seconds. The processor was averaging 67 percent busy and total paging was 33 pages per second with the CICS regions bearing the brunt of the paging activity. DASD channel utilization was in the 17 to 20 percent range and there were no DASD volumes with excessive response times.

The main problem with the I/O subsystem was cross system contention. A set of summary statistics for the CICS regions and the system was developed and tracked over the life of the project. To provide a consistent picture of the user workload MRO, routing transactions *are not* included in the reported transaction volumes. At the beginning of the project these statistics were:

Region	Tran Per/Min	Avg Resp
Tow Truck	250	.33
Data Entry	109	.34
Dist Office	476	1.32
Nov 86 — All	835	.90

Date	%CPU Busy	Page Rate	Tran Rate	Avg Resp	Comment
Nov 86	67	33	835	.90	

Tuning Project Goal

Overall the CICS workload was providing tolerable service to the user community. The response time objectives for the Tow Truck and Data Entry applications were well within the (informal) service level objectives established for those workloads. The (informal) service level objective established for the district office workload was .75 seconds average host response time; the actual response time was 1.32.

The primary motivation for the tuning project was the phased implementation of a new insurance claims system in the district office region and the addition of two dispatch offices to the tow truck system. It was felt that adding these new workloads to the existing system with no changes would deteriorate service to unacceptable levels.

The goal of the tuning project was to bring District Office response times within the service level objectives and accommodate the additional planned growth.

Tuning Strategy

A top down CICS tuning strategy developed by Eric Emanual (Candle Corp., Los Angeles, CA) was used to analyze the situation to find opportunities for improvement. This technique directs the analyst to begin with external system factors to ensure the CICS workload has the resources needed to process its workload. After the external resource shortages and bottlenecks have been corrected, then the tuning efforts should be directed toward the internal factors within CICS to improve performance and resource consumption. The primary focus of the tuning project was on the external factors affecting performance.

Initial Assessment

The primary performance bottleneck was a shortage of real and virtual storage. The paging rate was 33 and the application regions were experiencing frequent program compressions. This, coupled with the MRO configuration that combined terminal definitions and workload in a single region, contributed to the elongated and erratic response times the district office workload was delivering. The I/O subsystem experienced occasional cross system contention that was devastating to response time and throughput when it occurred. This was not considered a major factor in the overall response times being delivered.

Initial priority was to reduce the paging rate by acquiring additional real storage. The second priority would be a reconfiguration of the CICS regions to better utilize the features and capabilities of MRO and reduce the virtual storage constraint. Once these two changes were implemented, the workload would be re-evaluated to determine additional tuning opportunities.

Memory Upgrade

Real storage was upgraded in January, 1987 from 24MB to 32MB with predictable results. One month after the upgrade, demand paging dropped from 22 to two pages per second and processor utilization increased to 73 percent. In systems in which there are a limited number of active tasks, a high paging rate can throttle CPU utilization because it is possible, and likely, that enough of the regions may be waiting for page faults that the processor will not have enough active tasks to keep it busy. This did not, however, significantly reduce response times. The overall response time was .74 seconds. In February the regions had the following usage profiles:

Region	Tran Per/Min	Avg Resp
Tow Truck	259	.26
Data Entry	116	.32
Dist Office	518	1.07
Feb 87 — All	893	.74

The summary statistics for the system after the memory upgrade were:

Date	%CPU Busy	Page Rate	Tran Rate	Avg Resp	Comment
Nov 86	67	33	835	.90	
Dec 86	66	23	819	.85	
Jan 87	70	9	869	.73	Memory Upgrade
Feb 87	73	2	893	.74	

Adding the real storage did not provide the response time improvement that was hoped for. While it did remove the paging problem, it put further emphasis on a bottleneck that needed to be dealt with, VSCR. Although there was adequate real storage, the CICS region was limited to 16MB of virtual addressability and the size of the workload was pushing this limit. To resolve the VSCR problem in the shortest time frame, the MRO feature would be used to further divide the workload.

Distinctive Software for CICS

Electronic Mail System (EMS)

- Host based Electronic Mail System
- Scheduling & Calendaring
- Task List Management
- Full Function Full Screen Editor
- 3270 & Line - Mode Interfaces
- Extensive Security
- Unlimited Mailboxes
- Route Lists
- Many More Features

Side-CICS the Pop-Up Utility

- Screen Print (w/BMS Pagesets)
- Quick Message Pad
- Calendar and Scheduling
- Terminal Scan/View
- Instant CEDF Activation
- Calculator (Hex and Decimal)
- Access to CICS Utilities
- Telephone Directory

Gateways to the World

- Western Union Easylink
- ITT Timetram
- Graphnet Freedom Network
- TRT Multispeed
- DDD Delivery
- Tymnet Outdial
- ASCII Line - Mode Passthru
- Custom Gateways Quoted on Request

CICS ABend - Catcher

- Capture Key Transaction Data
- Compute Offset of Failure
- Save Actual End - User Screen
- Soft - Land User
- Immediate Notification of ABend
- TSO, CICS, and EMS Notification
- Batch Reporting Facility
- Control Dump Dataset Usage

CAS

Computer Application Services, Inc.
15560 Rockfield Boulevard • B2
Irvine, California 92718
(714) 859 - 2274 • Telex 755741

FETCH DESTROYS CICS STRESS

FETCH is a powerful solution to the problems that cause CICS stress. FETCH will dramatically improve CICS performance from day one. Just how good is it, really? Well, you could say FETCH is dynamite. Install FETCH on your CICS and get instant relief. Relief from CICS "bottleneck" problems like virtual storage constraints and slow response time.

You see, FETCH goes right to the heart of the matter by satisfying an unlimited number of CICS load requests simultaneously. So FETCH improves your response time right off the bat. Also, FETCH eliminates the need to define high use application programs resident for performance purposes. FETCH's unique multi-thread load mechanism will load programs as quickly as if they were core resident. And, if you're an XA shop, our FETCH/XA product will let you take advantage of address space above the XA line without any program changes.

- Reduce resident program storage requirements by 90% or more.
- Eliminate 87% of program wait on the load library and instantly improve response time.
- Reduce I/O at least 50%.
- Reduce compression problems.

FETCH™ & FETCH/XA™
OPERATES UNDER MVS, SP & XA

AXIOS PRODUCTS, INC.
1455 Veterans Highway
Hauppauge, NY 11788
(516) 348-1900

FREE 30 DAY TRIAL

CIRCLE #84 on Reader Service Card ▲

MRO Reconfiguration

The typical symptom of a virtual storage constraint problem in CICS is frequent program compressions and this was the case with both of the district office application regions. The addressability requirements for the application program code and file buffers routinely exceeded the size of the region that was limited to 16MB due to the level of CICS software. When this situation occurs, CICS performs a program compression to remove non-active programs that provide space for additional programs as they are called in for execution. Deleting and reloading program code on a regular basis is a high overhead task. If the workload within the region can be split and put in two regions, there would be double the virtual addressability for program code and file buffers, reducing the over commitment of virtual addressability and eliminating the need to do frequent program compressions. However this change will increase the real storage requirements of the workload. The existing CICS configuration was reviewed to look for opportunities to exploit this feature. Initially the CICS workload had the configuration in Figure 1.

FIGURE 1

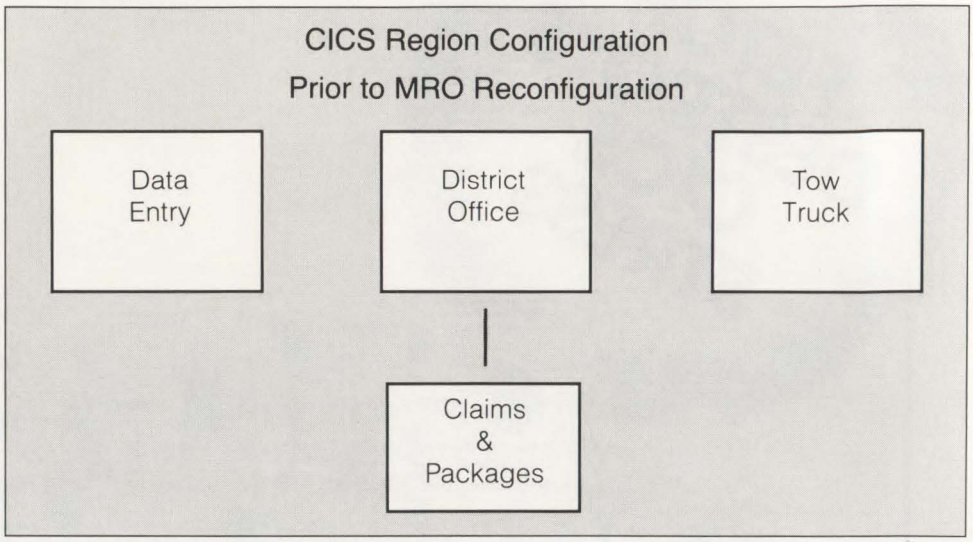

The MRO reconfiguration was completed in two phases. Initially the purchased packages were moved into a separate application region. The second phase involved moving all applications out of the terminal definition region making it a traditional terminal owning region. This phase involved creating an additional application region to contain District Office applications.

When the additional regions were created, the SRM parameters for CICS were reviewed and adjusted to give the terminal owning region the highest dispatching priority and all other application regions were placed in a single mean time to wait dispatching priority. The regions were grouped in this manner to maximize traffic to the I/O subsystem. There were some other adjustments made to the SRM parameters, such as storage isolation, but the impact of these changes was hard to determine since it was done in conjunction with the implementation of the new MRO configuration. At the conclusion of the MRO reconfiguration, the CICS workload had the configuration in Figure 2.

The MRO configuration was the most significant change implemented during the project. The additional regions drove the processor utilization up to 80 percent and increased demand paging to eight pages per second. This higher processor utilization translated into 34 percent reduction in the response time of the District Office applications. There were additional benefits outside the performance area from this change that are reviewed in the summary section of the article. The month after the reconfiguration was implemented the regions had the following usage profiles:

Region	Tran Per/Min	Avg Resp
Tow Truck	249	.33
Data Entry	103	.28
Dist Office	510	.71
Apr 87 — All	862	.55

Summary statistics for the system were:

Date	%CPU Busy	Page Rate	Tran Rate	Avg Resp	Comment
Nov 86	67	33	835	.90	
Dec 86	66	23	819	.85	
Jan 87	70	9	869	.73	Memory Upgrade
Feb 87	73	2	893	.74	
Mar 87	76	5	893	.47	MRO Reconfiguration
Apr 87	80	8	862	.55	

Second Assessment

At this point it was felt that the software configuration would benefit from a processor upgrade and support the additional workload planned for the Claims and Tow Truck system. Response time for the District Office applications was down to .71 seconds, .04 below the service level objective for this workload. The bottleneck was now the processor; average prime shift utilization was now up to 80 percent. Steve Hackenberg's article in MAINFRAME JOURNAL, "A Heuristic Modeling Approach to CICS Capacity Planning" (May/June 1988), presents a software queuing model that shows CICS response time will have a significant wait due to dispatch delay in a system with this level of processor utilization. Upgrading from a 3081-GX to a 3081-KX would be the next phase in the tuning project. After the processor upgrade, a CICS software upgrade was planned to provide additional VSCR. A MVS software upgrade was planned for software currency purposes.

Processor Upgrade

Upgrading the 3081 processor from a GX to a KX provided a 40 percent increase in MIPS from 11 to 14. This upgrade took place during the month of May. The 40 percent increase in processor speed resulted in a 42 percent reduction in response time and a 14 percent drop in processor utilization. It was concluded that this change significantly reduced dispatching wait as a component of CICS response time. This result is consistent with Hackenberg's software queuing model. Increasing the speed of the processor also has the effect of driving the I/O subsystem harder. VSAM buffer usage would be the next component of the environment that would need to be addressed. In June the regions had the following usage profiles:

Region	Tran Per/Min	Avg Resp
Tow Truck	249	.23
Data Entry	111	.24
Dist Office	564	.37
Jun 87 — All	924	.32

BEFORE

your systems support becomes terminal...contact ISS.

Some of our services include:

- System generation DOS/VSE, MVS VM, CICS, IMS, VTAM
- Conversion CPU upgrades: 43xx, 303x, 308x, 309x
  Telecommunications: BTAM, VTAM
  OEM Software Packages: Disk and Tape Management
- Performance evaluation & tuning
- Resource capacity planning
- On-demand technical support
- 24-hour emergency support
- Site survey (software & hardware)
- On going systems support
- Remote systems support

USA (818) 966-0227
LONDON 01-642-4242
BRUSSELS 01-736-9713

A Galliard Company

ISS

CIRCLE #94 on Reader Service Card ▲

Summary statistics for the system were:

Date	%CPU Busy	Page Rate	Tran Rate	Avg Resp	Comment
Nov 86	67	33	835	.90	
Dec 86	66	23	819	.85	
Jan 87	70	9	869	.73	Memory Upgrade
Feb 87	73	2	893	.74	
Mar 87	76	5	893	.47	MRO Reconfiguration
Apr 87	80	8	862	.55	
May 87	69	7	887	.46	Processor Upgrade
Jun 87	66	13	924	.32	CICS X/A Upgrade

CICS Software Upgrade

CICS Version 1.6.1 with the X/A feature would place the VSAM file buffers above the 16MB line. This change provided substantial Virtual Storage Constraint Relief since file buffers can account for up to 50 percent of the addressability required for user applications. However the additional virtual addressability did carry a real storage price tag. The production regions were upgraded to this version of CICS with no changes in the file buffer specifications and the working sets of the production regions increased to the point that overall paging increased from seven to 11 pages per second. Once the system impact of the software upgrade was established, VSAM file buffers were increased to keep the index set records of active VSAM datasets memory resident and string numbers were increased to eliminate significant string waits.

These changes did not have a significant or major impact on the overall statistics for the regions or the system. However it did tend to level out or flatten the peak period response time spikes that were occurring. This change completely eliminated program compressions during the normal operation of the system. The below-the-line addressability that the VSAM file buffers released was more than enough to store all programs.

While reviewing the VSAM buffer allocations, the use of Locally Shared Resource (LSR) buffer pools was investigated as a possibility. This plan was dropped because it was not possible to allocate resources needed to test it or check application code to ensure the LSR feature could be implemented without causing the problems associated with it.

In July the regions had the following usage profiles:

Region	Tran Per/Min	Avg Resp
Tow Truck	260	.22
Data Entry	103	.26
Dist Office	567	.37
Jul 87 — All	930	.32

Summary statistics for the system were:

Date	%CPU Busy	Page Rate	Tran Rate	Avg Resp	Comment
Nov 86	67	33	835	.90	
Dec 86	66	23	819	.85	
Jan 87	70	9	869	.73	Memory Upgrade
Feb 87	73	2	893	.74	
Mar 87	76	5	893	.47	MRO Reconfiguration
Apr 87	80	8	862	.55	
May 87	69	7	887	.46	Processor Upgrade
Jun 87	66	13	924	.32	CICS X/A Upgrade
Jul 87	69	12	930	.32	

Operating System Upgrade

No tuning changes were implemented in the three months following the CICS X/A software upgrade. During this period the transaction volume and all key system usage indicators held constant. The next major change to the system was an MVS/XA operating system upgrade to version 2.1.7 from version 2.1.3. The primary reason for this change was software currency. This upgrade was implemented during the month of October, 1987 and it doubled the paging rate in the system. One month after the upgrade the paging rate was 25 pages per second and processor utilization was up to 73 percent. The key point here is that response time was still in the .3 second range in spite of the high resource utilization rates. In December, 1986 the paging rate was 23, processor utilization was 66 percent and response time was .85 seconds. It was concluded that the additional address spaces created by the MRO reconfiguration account for the ability to tolerate higher paging rates because the impact of a page fault on the overall workload is reduced when there are other address spaces available for dispatch. Creating additional workloads made it possible to more fully utilize both CPs of the 3081. These conclusions are also consistent with Hackenberg's software queuing model. The summary statistics at the end of the project were:

Region	Tran Per/Min	Avg Resp
Tow Truck	285	.31
Data Entry	92	.31
Dist Office	609	.41
Nov 87 — All	986	.37

Summary statistics for the system were:

Date	%CPU Busy	Page Rate	Tran Rate	Avg Resp	Comment
Nov 86	67	33	835	.90	
Dec 86	66	23	819	.85	
Jan 87	70	9	869	.73	Memory Upgrade
Feb 87	73	2	893	.74	
Mar 87	76	5	893	.47	MRO Reconfiguration
Apr 87	80	8	862	.55	
May 87	69	7	887	.46	Processor Upgrade
Jun 87	66	13	924	.32	CICS X/A Upgrade
Jul 87	69	12	930	.32	
Aug 87	69	13	954	.35	
Sep 87	70	13	961	.30	
Oct 87	68	19	1007	.33	MVS2.27 Upgrade
Nov 87	73	25	986	.37	

Summary and Conclusions

During the 12-month period reviewed here, the overall CICS transaction volume increased by 18 percent and the response time decreased by 60 percent. It took a combination of hardware upgrades and software adjustments to make this hap-

See CICS *page 107*

FIGURE 2

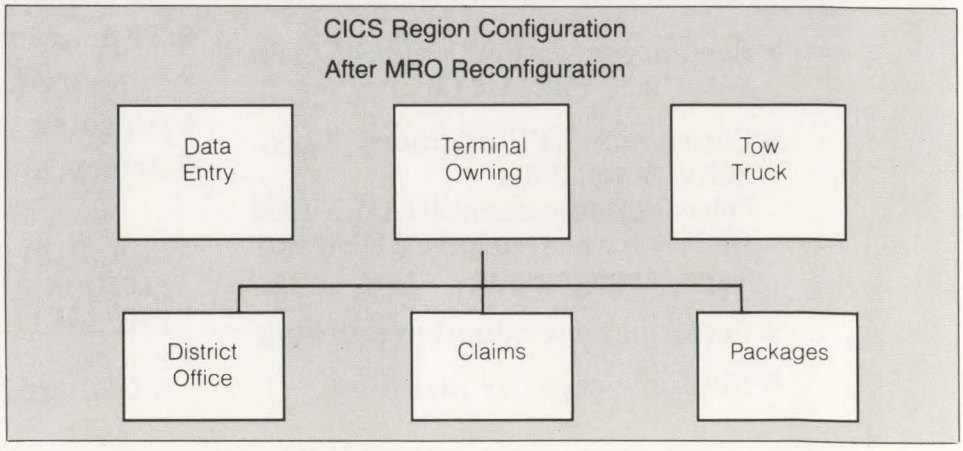

Without The VITAL SIGNS™ VM Performance Monitor, You Can Relate CICS Performance To VM Resource Usage For Improved Service and Planning...

But It *Isn't* Easy!

Now you can further improve service to end users and better plan growth by combining VM and CICS performance data. Quickly and accurately. When VITAL SIGNS™ is used in conjunction with Landmark Systems' The Monitor for CICS.™

VITAL SIGNS is the *first and only* VM performance monitor to include an interface to a CICS performance monitor. Through the interface, CICS performance data from The Monitor for CICS can be combined, by date and time, with VM data in the VITAL SIGNS performance database. Using the VITAL SIGNS Report Writer and Plot Generator, CICS performance information (e.g., response time) can be accurately associated with coinciding VM performance data (e.g., total system CPU utilization) to provide true correlations for identifying performance problems and trends. No more manual comparisons. No more guessing. And that's just the beginning!

VITAL SIGNS is also the *first and only* VM performance monitor to include DASD modeling. It allows you to perform "WHAT IF" analyses for such things as verifying the need for DASD upgrades, determining the effect of different file block sizes, and much more.

Of course, VITAL SIGNS also has both real-time and historical reporting facilities fully integrated into a single comprehensive product sold on a site license basis. You can have all major performance measures in both tabular and color graphics formats. Comprehensive DASD seek analysis facilities, such as a Volume Seek Map, are included. Exception reporting facilities provide early warnings of approaching problems. And an easy to use data extraction facility lets you pull specific data from VITAL SIGNS for further analysis by popular statistical or spreadsheet products on PCs or mainframes.

To find out first hand why VITAL SIGNS is more than just a VM performance monitor, call us today at ***800-826-0313*** (***612-542-1072*** in Canada and Minnesota).

BlueLine™
SOFTWARE INC.

New Advances in Performance, Productivity and Planning.

CIRCLE #48 on Reader Service Card ▲

1500 South Lilac Drive, Suite 340
Minneapolis, MN 55416
800-826-0313
612-542-1072 (In Canada and Minnesota)

EMC's ORION:
The Cost-effective Solution to Your I/O Performance Problems.

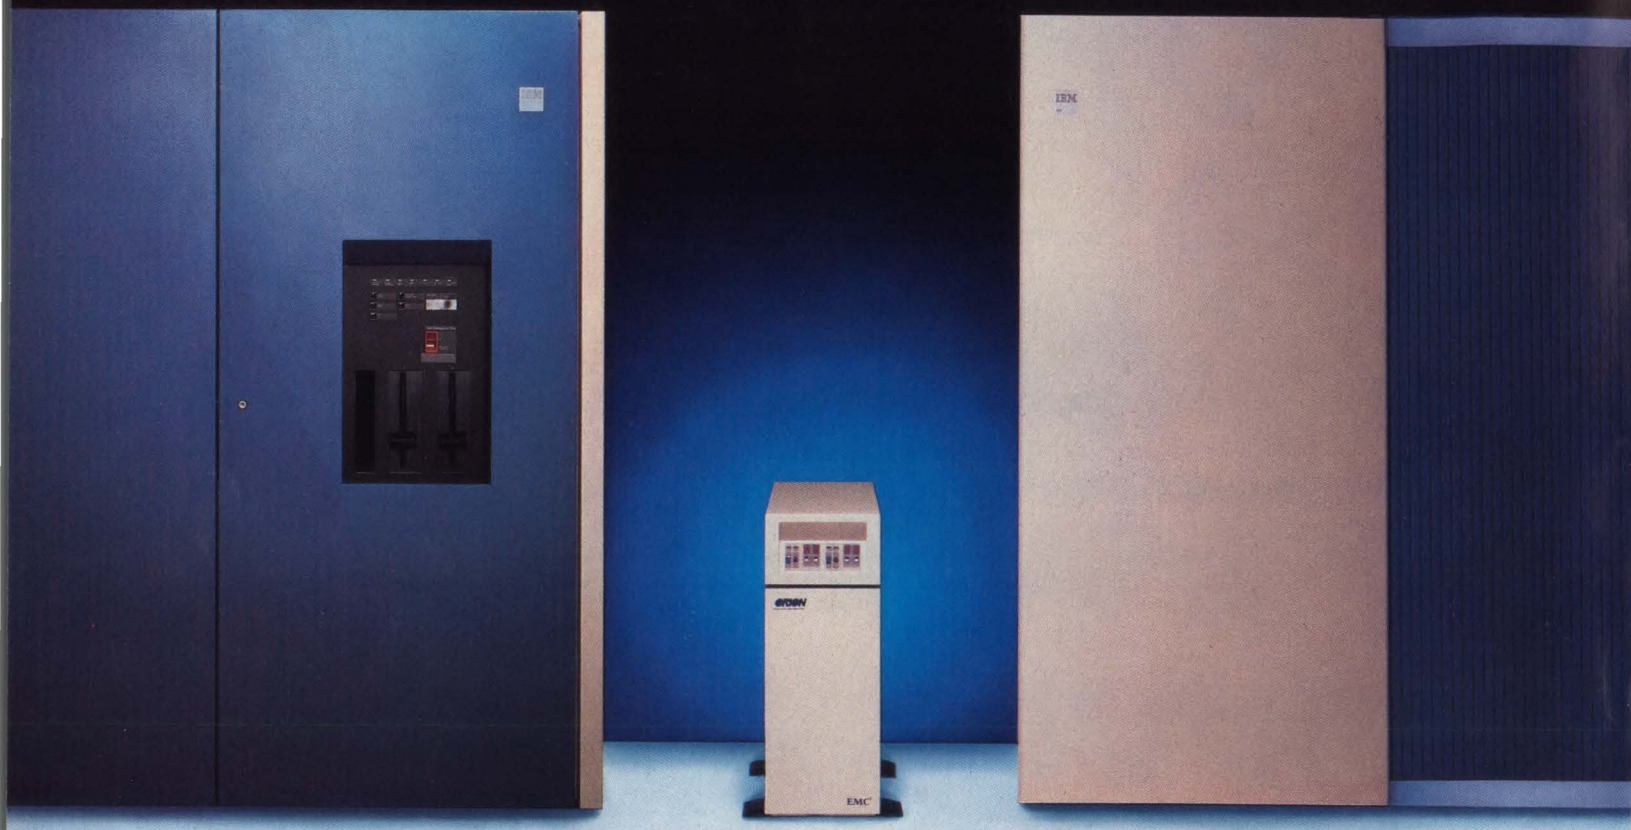

EMC's ORION Solid State Disk Subsystem

Maximum Performance

EMC's ORION is the fastest solid state disk subsystem available for your mainframe computer. ORION features a technologically advanced design enabling you to receive unprecedented performance gains.

An integrated 3880 type storage director, in addition to features inherent in solid state technology, gives ORION an access time of 0.1 millisecond — a performance milestone.

What's more, EMC protects your investment by making ORION compatible with all IBM 370 and IBM PCM computers. Therefore, performance boosts are realized well into the future, when CPU upgrades become necessary.

EMC's ORION — the I/O solution you have been waiting for to maximize your system's performance and productivity.

ORION: High Performance/Low Cost

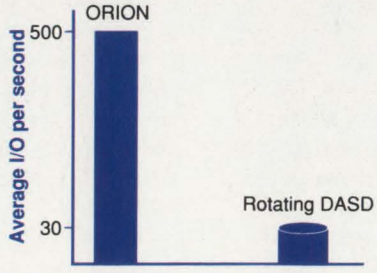

Minimum Price

EMC's ORION Solid State Disk Subsystem is the lowest cost solid state solution to your I/O performance problems. Its state-of-the-art design gives you a substantial return on your investment for years to come.

ORION's small footprint and low power requirements eliminate costly computer room renovations. Ease of installation and low cost of ownership make ORION an expedient and economical solution to your I/O performance problems.

EMC's ORION — the most economical performance boost for your current and future mainframe systems.

EMC²
The System Enhancement Company.

IBM is a registered trademark of International Business Machines Corp.

For more information, call:
1-800-222-EMC2
In Mass., call (508) 435-1000

Copyright 1988 EMC Corporation

CIRCLE #95 on Reader Service Card ▲

ISPF SPELLS PRODUCTIVITY

- Flexible
- Powerful
- Productive
- Menu-Driven
- Context-Sensitive Help
- Fill-in-the-Blank Screens
- Intelligent, Personalized Defaults
- Designed with the Programmer in Mind

By Jon E. Pearkins

Not phrases you would have associated with IBM software just a few years ago. In fact, the major criticism of the IBM mainframe environment has always been the labor-intensive nature of applications programming and technical support. And much of the blame has been attributed to the complexity of systems software.

Interactive System Productivity Facility/Program Development Facility (ISPF/PDF) has come a long way to change all that. Originally introduced at the end of the 1970s, ISPF/PDF has grown to the point where most utility-type functions can be done on-line without coding JCL or remembering the format of TSO commands. And many software products from both IBM and other vendors provide a set of panels for use on-line. ISPF provides the flexibility to add these panels as new items on any existing menu.

What is ISPF?

So, what is ISPF? IBM defines it as "a dialog manager that provides application developers with services to shortcut the coding required to produce interactive applications." Then, what is ISPF/PDF? "A program operating under ISPF that is used to quickly and economically develop and maintain applications — batch and interactive. It gives programmers convenient access to a set of tools used in all facets of on-line program development and testing." Taken together, ISPF gives you an excellent full-screen editor and productive set of interactive panels for executing utility functions: delete, define, copy, rename and list information about datasets. In short, ISPF is an editor and on-line DASD utility.

You can also develop applications that use ISPF to provide screen handling and other routine functions, but that is another

FIGURE 1

ISPF 3.4 — The Most Productive Panel

DATASET LIST UTILITY

```
OPTION = = = >

  blank    — Display dataset list*          P    — Print dataset list
  V        — Display VTOC information only  PV   — Print VTOC information only

Enter one or both of the parameters below:
  DSNAME LEVEL          = = = > YOURID
  VOLUME                = = = >

SPECIFY THE FOLLOWING, IF DISPLAYING A LIST OF DATASETS:
  DISPLAY FORMAT OPTION     = = = > QUICK    (QUICK, SHORT, or LONG)
  CONFIRM DELETE REQUEST    = = = > YES      (YES or NO)

*The following line commands will be available when the list is displayed

  B — Browse dataset                  C — Catalog dataset
  E — Edit dataset                    U — Uncatalog dataset
  D — Delete entire dataset           P — Print entire dataset
  R — Rename entire dataset           X — Print index listing
  I — Dataset information             M — Display member list
  S — Dataset information (short)     Z — Compress dataset
```

story. In this two-part article, we will explore productive ways of using ISPF that many are not aware of.

VM, MVS and VSE

ISPF is available in MVS under TSO, VM under CMS and VSE under ICCF.

ISPF under VM provides almost all of the features of the TSO implementation of ISPF. Although the ISPF Editor seems to have borrowed heavily from XEDIT's good points while improving on the bad, you can choose to use XEDIT instead of ISPF's Editor under VM. Unlike the TSO implementation, if you edit a file in VM, CMS will automatically create it if it does not already exist. However the ISPF Editor's default format for automatically-created files is fixed length (records of 80 bytes) while XEDIT defaults to variable-length records.

The ISPF Library is the most functional dataset type of ISPF. You have a choice of three implementations in VM:

1. TXTLIB
2. MACLIB
3. Standard CMS files.

TXTLIB and MACLIB are similar to MVS' Partitioned Dataset (PDS) in which individual members are stored in a single dataset. However in the standard CMS files implementation, each member is stored in a file with the Library's directory and member statistics stored in another file having a strange-looking name. But no matter which of the three methods you choose, the ISPF Library provides an excellent way to organize files into groups without resorting to multiple minidisks.

ISPF does need virtual memory to function. Unless you are editing extremely large files (more than 10,000 lines) or executing IBM's INFO products from within ISPF, 1.5MB will handle even the most complex task in an MVS/XA environment assuming the Link Pack Area (LPA) is being exploited. For MVS/SP, two megabytes would be required. Two megabytes would also be the appropriate CMS virtual machine size in VM as well as the size of the ICCF interactive partition used to run ISPF in VSE.

FIGURE 2

Using Split Screens

```
PF2   — create or move the second screen:
        cursor position = top
PF9   — make the other (split) screen the
        dominant one
PF12  — move the cursor to the command
        line of the other screen
=X    — exit split screen mode, removing
        the current screen
```

The Most Productive

Assuming your installation has sensible dataset naming conventions, you will probably find ISPF menu item 3.4, the Dataset List Utility, the most productive panel (see Figure 1). From it you can display a list of all datasets with the same high-level name and, optionally, other common characters found anywhere in the dataset name. The list is displayed in alphabetical order by dataset name. However you can SORT the list by DASD volume serial number, size (largest first), expiration date, last referenced date, creation date or percent used.

Once this list is displayed you can perform common actions on the datasets: edit, browse, delete, rename, print and display allocation information. Assuming the dataset is a PDS, as all ISPF libraries are in MVS, you can also compress it, display or print a list of members and then perform these same functions on the ISPF Library members themselves.

When you need to do some other dataset functions such as allocate a new dataset, copy or move an existing dataset or look at the output from a batch job, you can invoke the ISPF split screen capability (PF2, by default) and access these functions through menu selections 3.2, 3.3 and 3.8 (see Figure 2).

There is one caution with Split Screen mode. If you are editing an input dataset in one screen and submitting a batch job or TSO CLIST that uses it from the other, do not forget to save the file before submitting the job that will use it.

If you are a long-time ISPF user and have not looked at 3.4 in the last few releases, give it a try. It can save you remembering and typing a lot of dataset names.

Help Function

Reading the manual is normally the best way to learn about features of IBM software that can help make you more productive. With ISPF, the Help function is the best way to learn. Help (PF1) is context-sensitive. If you are editing a file, Help tells you about editing. However if you are not interested in the current context, hit PF1 again and you will see the table of contents for all available Help. If you want to use the Help panels as a mini-course, they are available as selection T (Tutorial) from the ISPF Main Menu.

If you have just received an error message and then hit PF1, you will get a more detailed explanation of the error condition sometimes with a suggestion to solve it. For example, when concluding an Edit, if you get IEC0321 E37-04 from TSO and then SYSTEM ABEND '0E37' from ISPF, hit PF1 twice and ISPF will tell you to split the screen and compress the PDS. You can do this without losing your changes.

Anywhere in Help, you can type the letter "I" to get into the Help Index. You

System-managed Storage

Sterling Software's Storage Management Strategy

1988 DMS/OS and VAM together give you system-managed storage today— automatic volume allocation, reporting, archival, backup, idle space release, PDS compression, DASD space billing, and dynamic DASD reconfiguration.

1989 DMS/OS and VAM will support ESA when it is released— along with all other MVS operating systems on any IBM/PCM mainframe.

1990 Sterling Software will begin releasing a new family of storage management tools called Corporate Storage Management (CSM), which will work beside and beyond IBM. In addition to functions in DMS/OS and VAM, CSM will give you intelligence, file compression, connectivity, performance and capacity management, and management-style reports.

Today and in the future, Sterling is the solution

STERLING SOFTWARE
Systems Software Marketing Division

Pushing the limit in system-managed storage

For a free copy of our position paper on Corporate Storage Management and the future of system-managed storage, call us at (916) 635-5535 or write today, 11050 White Rock Rd, #100, Rancho Cordova, CA 95670.

CIRCLE #73 on Reader Service Card ▲

FIGURE 3

Line Command Tips

RRn — Repeat a block of lines
TS — Split a line at the cursor
O — Combine two lines
OO — repeat characters over many lines
XX — do not display a block of lines
))n — shift right "n" columns
.label — label a location in a file
TABS — display/modify tab positions
COLS — display column numbers

can then type the first letter of the topic you are interested in and see a list of all topics beginning with that letter. If you see one you like, type the letter-number code for the particular topic you want and it will be displayed.

PF3 gets you all the way out of Help, but what can you do to move up one level of Help? PF7 (UP), normally used to Scroll Back in a file or list, will take you back to the previous Help Menu. PF10 (LEFT) provides the normal Scroll Back function by displaying the previous page of Help. To display the next page of Help, just hit the ENTER key.

Line Commands

Line Commands are those ISPF Editor codes that you type over the sequence number for the (first) line involved (see Figure 3).

Using the C (Copy) Editor Line Command to make several copies of a block of lines can be time-consuming. For each copy, you must mark the block of lines with CCs on the first line and the last line and place an A (After) or B (Before) on the line where they will be inserted. If all of the copies are going in the same location, use the R (Repeat) command. Mark the beginning of the block of lines with an RRn in which "n" is the number of copies you want and the end of the block with RR. The result: "n" copies of the block of lines will be inserted immediately after the last line of the block.

Inserting something into the middle of a line can be a frustrating task. If there is enough room on the line, enter the NULLS Primary Command. From then on the INSERT key on the 3270 will allow you to type characters without retyping the rest of the line. But what do you do when the line gets full? The TS Line Command will break the line into two lines using the position of the cursor to determine where the line should be broken. If you find yourself using TS a lot, set one of the PF keys you do not use (=0.3) to ":TS", position the cursor where you want to split a line and hit the PF key.

Conversely, to combine two lines, use the M (Move) and O (Overlay) Line Commands. For example, if the first line ends at column 59 and the text on the second line should immediately follow, the first line should be blank *after* column 59 and the second blank *before* column 60. One line would be marked with an M and the other with an O.

A similar technique may be used to add the same information to multiple lines. For example, suppose that you have 20 consecutive lines where you want to add an X in column 72. To do this, insert a new line anywhere in the file containing just an X in column 72. Then use the M Line Command on this new line and mark the beginning and the end of the 20-line block with OO (Overlay Block).

If you have ever tried to look at two parts of the *same* file on the same screen using Split Screen, you are familiar with the ISPF "MEMBER IN USE" message. A much better alternative is the X (eXclude) Line Command. Mark the block of lines *between* the two parts of the file with XX on the first and last lines. The block will no longer be displayed, but its existence will be indicated by a dashed line and an "n LINE(S) NOT DISPLAYED" message. When you want to display the block again, the Show Line Command, S9999, can be keyed in the prefix area of the dashed line.

Shifting text right or left on one or more lines is a common requirement especially in structured programming in which indenting is the key to readability. Mark the text with))n or ((n at the beginning of the block of text and)) or ((at the end of the block: "n" indicates the number of columns to shift the block of text. The direction in which the parentheses point indicates the direction of the shift. For example,))37 indicates a right shift of 37 columns.

Primary Commands

Primary Commands are those ISPF Editor commands that are typed on the COMMAND ===> line near the top of the screen (see Figure 4).

Changing all occurrences of a character sequence (but only in part of a file) can even be slow when using PF5 (RFIND) and PF6 (RCHANGE) if the response time is poor and there are a lot of occurrences to change. Instead, mark the block of lines you want to change with XX on the first and last lines of that part of the file. The Primary Command CHANGE "string" "string" ALL X will change all occurrences in the block.

To redisplay the excluded lines that the XXs have created, the RESET Primary Command is probably the easiest. As well as displaying excluded lines, RESET cleans up line flags ("==CHG>" and "==ERR>") and removes special lines like column numbers and profile displays.

If you are typing a lot of Primary Commands that only change slightly from one command to the next, try preceding the first one by an ampersand ("&"). Instead of erasing what you typed on the Command Line after executing it, ISPF leaves it there for your modification.

If you get called away on a problem, by the time you get back to your desk you will have been automatically logged off. You can save yourself from losing all the changes you made by turning RECOVERY ON. Type the Primary Command

FIGURE 4

Primary Command Tips

RESET — display excluded lines; clean-up line flags
& — do not erase the command that follows
RECOVERY ON — protect Editor changes from systems crashes and so on
FIND P"#" — pattern matching
DISPLAY CC — show the printer carriage control for each line
SORT — sort a file by column contents
EDIT — edit a second file while in the first
LOCATE — find a label in Edit/Browse, area of interest otherwise
TABS — define the character for Logical Tabbing
HEX ON — display lines in both hexadecimal and character codes
FIND X"IF" — search for an unprintable character
TSO or CMS — execute a TSO or VM/CMS command within ISPF
PRINT-HI — print the physical screen
PRINTLHI — print all of a split screen when only partly shown

System 38, 4300, 9370 Users:

When it comes to price/performance of tape drives, only one system stacks up.

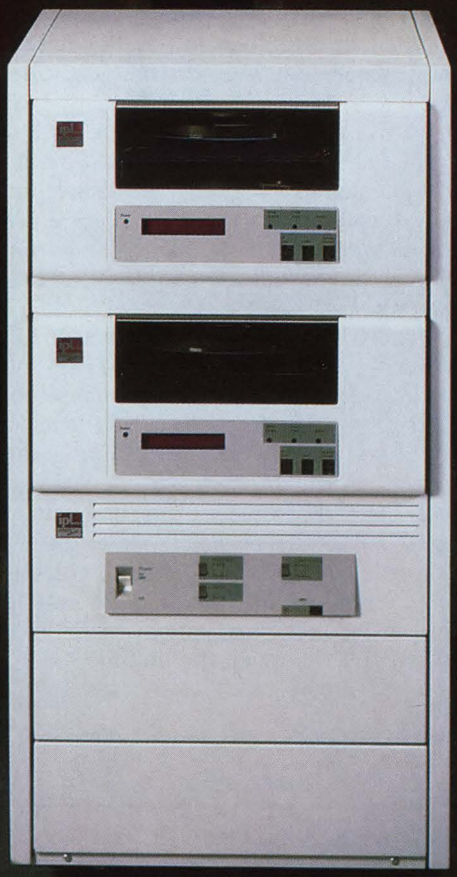

Introducing the 6800 Series tape subsystem from IPL. The 6800 provides a new level of tape drive performance for mid-range IBM processors. And its small size, low power and cool, quiet operation means the 6800 is ideal for office environments.

Our new rack mounted subsystem features an IPL intelligent controller, one or more tape drives, and room for future upgrades. Compared to other subsystems, here's how the 6800 stacks up. For backing up your system, it operates twice as fast as an IBM 3430. Yet it costs considerably less.

If you need more storage capacity, simply install a second tape drive in the rack. You'll have more storage capacity than an IBM 3422.

Now, for the optimum tape subsystem for your office, you have only three letters to remember: IPL.

1-800-338-8ipl

(in MA 617-890-6620)

ipl systems, inc.

360 Second Avenue, Waltham, MA 02154

IBM is a registered trademark of International Business Machines Corporation

CIRCLE #96 on Reader Service Card ▲

and ISPF will remember all of your changes whenever you are editing this file. RECOVERY is just one of many switches that the ISPF Editor remembers in its profiles. Profiles created for one dataset are applied to all others with the same lowest level name (or *filetype*, in VM).

Since lower-case letters are invalid in JCL and unacceptable to many other utilities and non-printable characters are invalid almost everywhere, Browse and Edit offer you ways to find them. The FIND Primary Command allows you to specify Picture Strings: FIND P''<'' for lower-case letters and FIND P'' '' for invalid characters. FIND P''-'' is useful for detecting any non-numeric characters in a data file in which one field is supposed to be numeric. Multicharacter Picture Strings are also useful. For example FIND P''@#@#'' would look for a letter followed by a digit followed by a letter followed by a digit with no intervening blanks. Blanks, letters and numbers can be included in Picture Strings when a combination of exact and pattern matching is required. For example, FIND P'' FOR###'' would search for a blank followed by the letters FOR followed by at least a three-digit number.

In Browse, you may not be seeing the Carriage Control (''CC'') in print files. Assuming the dataset is correctly specified with a DCB RECFM of A or M, the DISPLAY NOCC and DISPLAY CC Primary Commands determine whether or not the Carriage Control is displayed at the beginning of each line.

You do not have to write JCL to use DFSORT or one of its competitors to sort a file. As long as you want to sort Text fields (not Packed Decimal) and the record length is reasonable, the SORT Primary Command will do the job. Just specify the starting column and length of up to five fields you want to sort by. For example, SORT 10 5 1 3 will sort by the text in columns 10 through 14. Where the text in columns 10 to 14 is identical for two or more lines of the file, columns one through three would be used to determine the order.

Finally, if you want to edit another file without ''committing'' the changes you have made to the file you are currently editing, you can always use Split Screen. But if you already have something in the second screen that you do not want to lose, the EDIT Primary Command will let you ''go down one level'' and edit another file before you are finished with this one. When you hit PF3 to finish editing the second file, you will be right back where you were with the first file.

In the next article, techniques to help make your use of ISPF more productive will be examined. These techniques will include: moving around fast in files and in ISPF itself, the value of sequence numbers, living with or without unprintable characters, never getting stuck, ISPF tabbing simplified and the question of ISPF performance will be addressed.

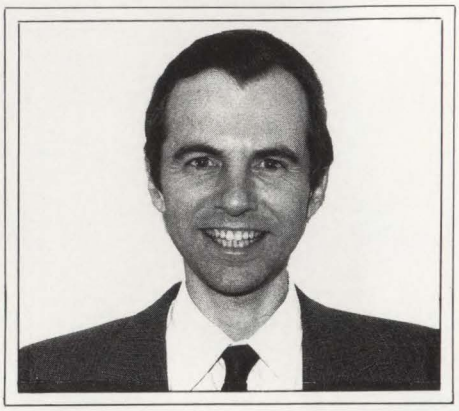

ABOUT THE AUTHOR

Jon E. Pearkins, CSP, CDP is president of Certified Software Specialists Ltd. He has spent almost all of his 15 years of data processing in technical support, the vast majority of which has been in an IBM mainframe environment. The company develops and supports its own line of systems software. Certified Software, 54015 Range Road 212, Ardrossan, Alberta, Canada, TOB OEO. (403) 998-0607, Fax 403-423-9995.

ISPF 2.3

The articles on ISPF in this and the next issue of MAINFRAME JOURNAL focus on ISPF Version 2, Release 2, most often in the MVS implementation. ISPF 2.3 has quite a number of improvements to offer. Here are some of the highlights that offer the most potential for productivity.

The Most Productive Panel, — Dataset List Utility (3.4), has many significant improvements:

- *A new Line Command, F, frees unused space in a dataset*
- *When Datasets are listed or Members of an ISPF Library are displayed, TSO Commands and CLISTs may be entered as well as the usual Line Commands. The Dataset or Member Name is added to the end of the command or inserted if you include a / in the command*
- *A new Line Command, =, repeats the CLIST, TSO or Line Command entered on the previous line*
- *E for Edit can be entered on the ISPF Library Member List displayed as a result of the M Line Command*
- *FIND and RFIND (PF5) can now be used to find any string in the Dataset Name field*
- *SORT and LOCATE can now be used on any field in the Dataset List with similar capabilities for Member Lists.*

Options 3.12, 3.13 and 3.14 have been added to ISPF to provide much needed comparison and multi-file search capabilities. Called the SuperC, SuperCE, Search-For and Extended Search-For utilities, they will be covered in future issues of this magazine once there has been an opportunity to gain some experience in their use.

We SELL-BUY-LEASE
EVERYTHING INCLUDING SOFTWARE!
If You're Looking for-
- QUALITY & SERVICE
- IMMEDIATE DELIVERY
- PROFESSIONAL EXPERTISE
- BEST PRICES ANYWHERE

R ASSOCIATES *PROVIDES IT ALL!*
"I GUARANTEE IT."
Herb Riley PRESIDENT

New or Pre-owned
4300's
30XX
PERIPHERALS
UPGRADES AVAILABLE
FOR FAR LESS THAN IBM CHARGES
ALL SYSTEMS GUARANTEED FOR IBM MAINTENANCE

YOUR OFFER COULD BE OUR PRICE
1-800-635-0909
R ASSOCIATES, INC.
1423 UPLAND HOUSTON, TX 77043

For CICS/VSAM Data Recovery, Integrity Solutions is the Global Choice.

These leading IBM mainframe companies trust the integrity of their corporate data to Integrity Solutions. They selected Integrity Solutions' DATA RECOVERY SYSTEM (DRS) because it provides the most comprehensive forward and backward recovery for files accessed by CICS/VS and VSAM batch processing.

Integrity Solutions introduced DRS more than 8 years ago, and with over 1000 product installations worldwide, Integrity Solutions will continue setting the data integrity standard. By providing leading-edge products and the highest level of support . . . . to serve your data recovery needs now and in the future.

Guided by what users say they need next, Integrity Solutions is setting the data integrity standard for the emerging electronic vaulting and mirror processing marketplace as well.

The DATA RECOVERY SYSTEM from Integrity Solutions, because your data is worth it.

Integrity
Solutions, Inc.

1-800-289-9900
1-303-794-5505

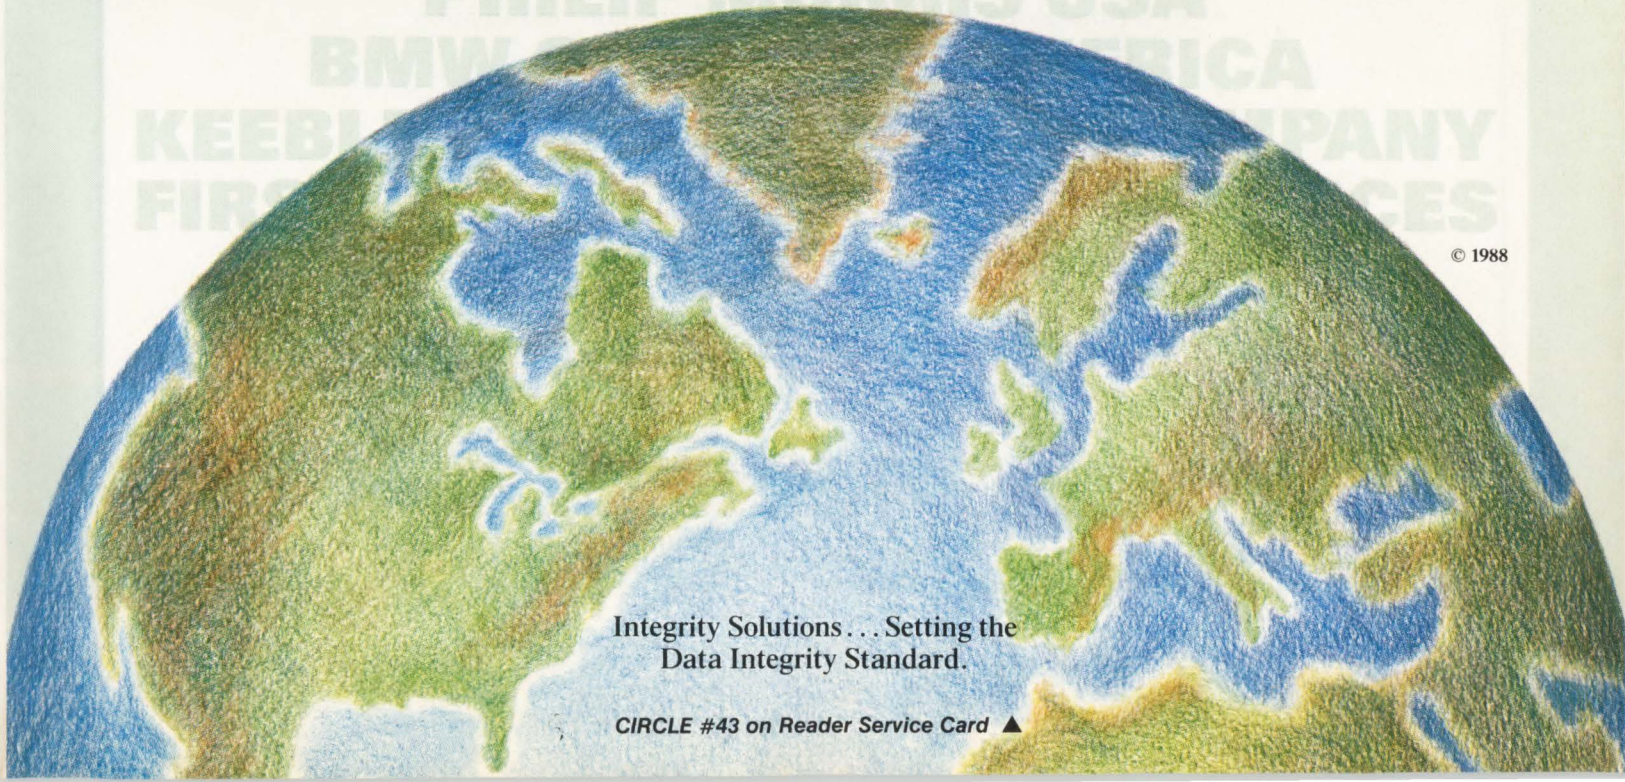

© 1988

Integrity Solutions . . . Setting the Data Integrity Standard.

CIRCLE #43 on Reader Service Card ▲

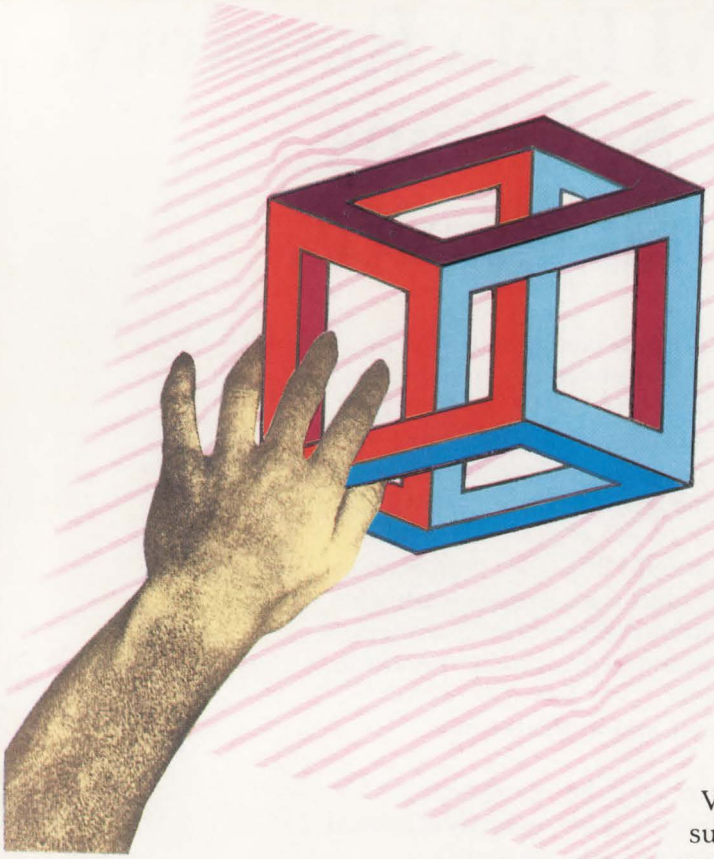

VPS Support Of Network Printers On MVS Is A Complex Task Made Simple.

VPS has been used to replace other products, such as: IBM's 328x/ADMPRINT/DSPrint, CICS supported printing, SASWTR®, RJE and many others, with a single task to drive all your 3270 family printers directly from the JES spool (including cross domain VTAM printers).

1400 Sites use VPS as their shop standard 3270 family printer driver.

- Printers supported include the full array of 3174/3274 attached printers, including IPDS support for 42x4's, HP and Xerox lasers, plotters and PC printers.
- VPS runs as a VTAM application. NO system modification. NO JES maintainance.
- Automatic forms control, full FCB support, dial-up PC printer, printer pooling and "hot" printers are all supported.
- Full screen "ISPF-like" command interface for CICS, TSO and ROSCOE permit end-user control of printers with totally menu driven command entry.

CALL or write for more information, or to arrange for a free TRIAL – ATTN: Marketing Dept.

LRS

Levi, Ray & Shoup, Inc.

Specializing in Computer Systems Software

2387 West Monroe • Springfield, Illinois 62704 • (217) 793-3800

Telex (510) 60-00675

® Trademark of SAS Institute, Inc., Cary, NC.

CIRCLE #121 on Reader Service Card ▲

The Basics of ACF/VTAM Logmode Tables

By Lloyd A. Hagemo, Jr.

If you are just beginning to use an SNA network, you know that there are many concepts to learn. One of these is the definition and implementation of the LOGMODE table. The misunderstanding of ACF/VTAM LOGMODE tables causes many problems in SNA networks. For the systems programmer who is new to SNA, it is sometimes difficult to understand how to define a LOGMODE table and when and where to use one. This article provides background information for using LOGMODE tables.

What is a LOGMODE Table?

A logon mode (LOGMODE) table contains one or more sets of session parameters that represent the session protocols that are to be used in an SNA session. The LOGMODE table is a set of ACF/VTAM macros that are assembled and link edited into the ACF/VTAM load library identified as VTAMLIB. ACF/VTAM loads the module into memory when the first major node that references the LOGMODE table is activated.

At some time after you define your initial LOGMODE table(s), you may find that you want to replace one. To replace a LOGMODE table, all of the major nodes that reference the table must be varied inactive at the same time. The new table must be assembled and link edited into the VTAMLIB and it must be available when the first major node that references the new table is activated. Contrary to popular belief, you do not have to shut down ACF/VTAM to implement a new LOGMODE table.

Two SNA network definition parameters are related to LOGMODE tables when defining resources: MODETAB and DLOGMOD. MODETAB points to the table that ACF/VTAM will use to build a logon request for an application. DLOGMOD points to the entry within the table that this resource will use.

The ACF/VTAM logon command allows you to override the DLOGMOD parameter, but not the MODETAB parameter. That is, you can override the table entry, but not the table. For any logon override to be successful, the DLOGMOD name must match an entry in the MODETAB table.

The default for MODETAB is ISTINCLM as defined in Appendix B of the *Advanced Communications Function for VTAM, Customization,* Version 3, SC23-0112. The default for DLOGMOD is IBM 3767, the first entry in the table. When you point a resource to a table and do not specify DLOGMODE, ACF/VTAM will use the first entry in the table as the default parameter.

All logon requests look up logon parameters in a table. If you have never needed to use or define a LOGMODE table, then either the default table or the subsystem being accessed is fulfilling your requirements.

The LOGMODE Table and Applications

The LOGMODE table does not have the final say in the session parameters. An application can redefine or override logon parameters. CICS, prior to release 1.7, is one example of a subsystem that overrides LOGMODE definitions. The CICS TCT definition governs which session parameters will be used. CICS has general SNA BIND images that it uses to talk to different device types depending on how the TCT definitions are specified. The parameters in the TCT definitions that are related to SNA session parameters are TRMTYPE, CHNASSY and BRACKET.

In CICS 1.7, the logon exit support has been improved to support more dynamic resource definitions. These definitions are generated based on the ACF/VTAM LOGMODE table values passed by ACF/VTAM to the CICS logon exit. CICS will dynamically build a TCT entry based on matching the BIND image with a predetermined resource pool. The LOGMODE table is used by ACF/VTAM to build this BIND image. Remember that an application can change the type of session parameters it wants to use. Be sure to check your applications documentation for any specific LOGMODE table requirements.

LOGMODE Table Macros

The LOGMODE table consists of three macros: MODETAB, MODEENT and MODEEND. MODETAB defines the start of the LOGMODE table and the CSECT name for the assembly. This is the name used as the MODETAB resource definition parameter. It should match the name used to link the assembled module into the VTAMLIB dataset. MODEENT associates a LOGMODE table entry name with a set of parameters that represent session protocols. MODEEND indicates the end of the LOGMODE table.

IBM supplies a default LOGMODE table with ACF/VTAM. You can modify or replace this table but it is a good idea to keep it intact. If you need to define different entries or other device support, you should do so in another LOGMODE table.

LOGMODE Table Sample Entry

Figure 1 illustrates one of the IBM default definitions in the sample table ISTINCLM for a 3278 Model 3 remote SNA terminal. Following the sample is a detailed description of each of the entries. Remember, an application can override any of these parameters at logon time.

FIGURE 1

Sample 3278 Model 3 Entry

```
D4C32783 MODEENT LOGMODE=D4C32783,
         FMPROF=03,
         TSPROF=03,
         PRIPROT=X'B1',
         SECPROT=X'90',
         COMPORT=X'3080',
         RUSIZE=X'87F8',
         PSERVIC=X'020000000000185020507F00'
```

The LOGMODE parameter is the entry name associated with this definition. This is the name that is coded on the DLOGMOD parameter in the resource definition and it is the table entry name.

FMPROF is the function management profile to be used for this session. The *SNA Reference Summary* contains the specific information about the support that is available with FM profile 03. FM profile 03 defines support for brackets and chains. The PRIPROT, SECPROT and COMPROT parameters further define which LU, primary or secondary, is going to be able to send an end bracket, win begin bracket contention situations and whether or not function management headers are used. In addition, the following data flow control RUs are supported with FM profile 03: CANCEL, SIGNAL, LUSTAT, CHASE, SHUTD, SHUTC, RSHUTD, BID and RTR. The application that the device is logged on to must also support this set of function-related services and protocols to communicate with this device.

TSPROF is similar to FMPROF, but instead of being function management related, it is used for transmission services. TSPROF 03 indicates that sequence numbers are used on normal flow RUs and that normal flow data is paced. The following transmission related RUs are supported: CLEAR, SDT, RQR, STSN and CRV.

PRIPROT defines how the application converses with the terminal. It further defines which FM and TS profile services are used. Hexadecimal 'B1' represented as X'B1' indicates a one-byte field containing Binary '1011' represented as B'1011' in the first four bits and B'0001' in the second four bits. The BIND RU contains the definition of this bit pattern in byte four found in the *SNA Reference Summary*. The first two bits indicate that multiple request chains are supported. The next two bits indicate that both definite and exception response modes are supported. B'0001' (the second set of bits) indicates that only the primary logical unit or application can send an end bracket.

SECPROT defines how the terminal converses with an application. This is defined in byte five of the BIND RU. Again, an X'90' translates into B'1001' and B'0000' in bits. The first bit indicates that the secondary supports multiple RU chains. You would expect this since the primary logical unit supports multiple request chains. The next bit indicates that the terminal supports only exception response mode. The terminal will not request definite responses from the application. The last four bits indicate that the terminal will not send an end bracket and does not support data compression.

COMPROT applies to both the primary and the secondary end, that is, application and terminal. This is defined in byte six of the BIND RU. The X'3080' definition translates into B'0011' B'0000' B'1000' B'0000'. The first bit indicates that brackets will be used and that the reset state is defined as being between brackets. The second bit in the first set indicates that bracket termination rule one is used. The next four bits are reserved and are all zeros. The next bit in the third set indicates half-duplex flip/flop mode. This means that an application will talk to the terminal and then send a Change Direction (CD) indicator. The CD indicator grants permission to the terminal to send data. The terminal sends its data followed by its own CD indicator. This applies only while the session is in the bracket state. The last set of bits indicates that the device is the first speaker and wins all bracket contention situations.

RUSIZE determines the maximum RU sent by the secondary and the primary logical units. The formula is M*2**N where, in the above definition, X'87F8' would translate to 8*2**7 or 1024 for maximum terminal RU send size and 15*2**8 or 3840 for a maximum application RU send size. IBM recommendations are usually acceptable for the device type and model number that you are defining.

PSERVIC is by far the toughest to define because it is device specific. It stands for presentation services and relates to session end-to-end communication. The *IBM 3270 Information Display System, 3274 Component Description and Programmers Guide*, GA23-0061, contains the description for this field illustrated in Figure 1. This description is found in the BIND RU discussion. The first byte indicates the session type. X'02' is defined for this device and means that it is an LU type two, 3270 device. The second byte indicates that this terminal supports the extended data stream commands and orders. The model 3270 sample definition shown in Figure 1 does not have this extended support.

The next field of interest includes the numbers X'1850' and X'2050'. X'1850' indicates that the terminal's default screen size is 24 by 80. X'2050' indicates an alternate screen size of 32 by 80. X'7F' indicates that the alternate screen is supported as defined by the X'2050' or 32 by 80. When using the PSERVIC definition, you must go to the component description manual for the device that you are defining to determine the contents of the field.

The sample in Figure 1 includes eight parameters. Besides these, there are five other parameters that could be part of a LOGMODE table entry. They are Class of Service (COS), Encryption (ENCR), PSNDPAC, SRCVPAC and SSNDPAC.

COS points to the class of service table to be used for the session. PSNDPAC, SRCVPAC and SSNDPAC are pacing parameters. PSNDPAC refers to primary send pacing, SRCVPAC refers to secondary receive pacing and SSNDPAC refers

to secondary send pacing. These three parameters override the resource definitions PACING and VPACING. These are performance-related items and are used to control the flow of data between the host and the device. ENCR specifies ACF/VTAM data encryption, a feature that is applicable only to MVS.

LOGMODE Table Usage

The LOGMODE table definition builds the CINIT RU used in session establishment. The CINIT RU has embedded within it the BIND image used to control the end-to-end protocol used for the session. An application's ACF/VTAM logon exit can modify the bind parameters prior to issuing the OPNDST macro that is used to establish sessions.

You can use an ACF/VTAM trace to see the CINIT and BIND RUs. Traces are a handy way to diagnose problems in SNA networks. The BIND RU is the first place to start in that process.

Conclusion

The *Systems Network Architecture Reference Summary,* GA27-3136, contains specific information on the SNA Transmission Services (TS) and Function Management (FM) profiles. The BIND RU is used in SNA to define session protocols. That RU can also be found in the *SNA Reference Summary.* Refer to that manual for a description of all the RUs identified in this article. If you want to start with some basics to help you learn SNA terminology, you might read the manual *Systems Network Architecture Concepts and Products,* GC30-3072, or *Advanced Communications Function for VTAM, General Information,* GC27-0608.

If you need to define a LOGMODE table for devices other than those defined in IBM's default table, an example is often helpful. The device's component description will usually contain either the format of the BIND RU or the actual LOGMODE table definition. IBM has a number of Washington System Center Technical Bulletins available. These contain helpful ACF/NCP and ACF/VTAM examples. SHARE and GUIDE are also good places to find resource people.

To understand the LOGMODE table, most people must do some background reading about SNA. I hope that this article helps you get started on the road to that understanding. As in all IBM definitions, there is nothing magical; it just takes time to get used to the definitions and how to use them.

ABOUT THE AUTHOR

Lloyd A. Hagemo, Jr. is manager of product development at BlueLine Software, Inc. He directs the development of the company's SNA products. Also, he has developed several VTAM applications and a product similar to NPDA. These applications are being marketed on DOS/VSE, MVS/SP and MVS/XA systems. His latest work is with GCS and ACF/VTAM running on VM. BlueLine Software, Inc., 1500 S. Lilac Dr., Suite 340, Minneapolis, MN 55416, (800) 826-0313.

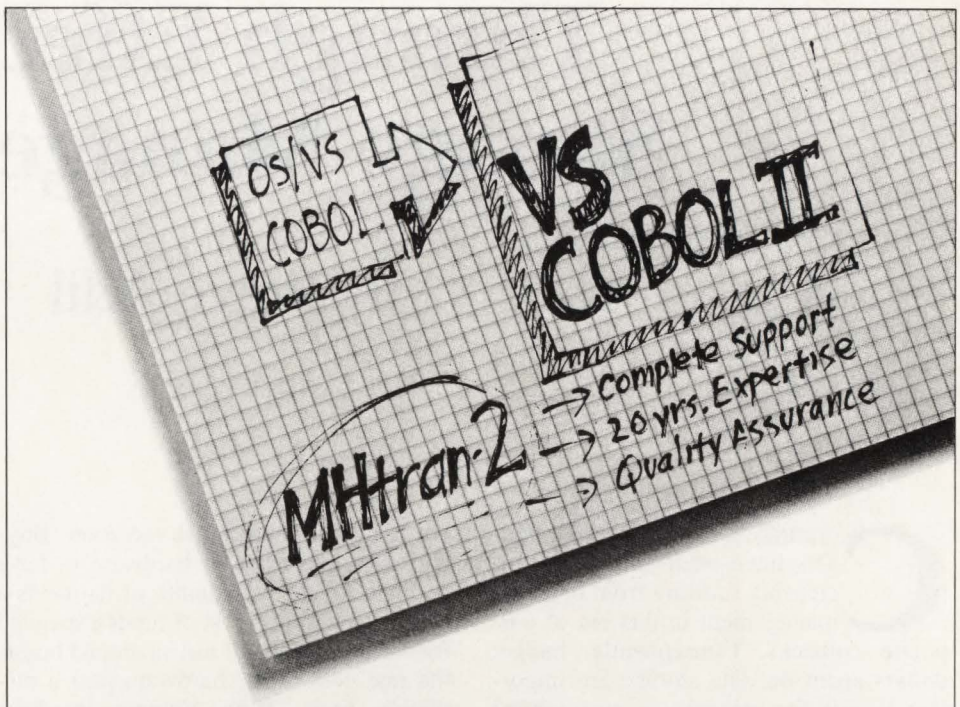

If you're planning the pivotal move to VS COBOL II...

Don't migrate alone.

MHtran-2 translates both batch and CICS programs from OS/VS COBOL to VS COBOL II. Completely. Accurately. Automatically.

And MHtran-2 guarantees you what no other translator can — the MHtran-2 migration experts who support your VS COBOL II migration from beginning to end. Every step of the way.

For information on a free 30-day trial, call 201-934-0022

MHtran-2

YOUR COBOL II TRANSLATION SOLUTION

PRINCE SOFTWARE PRODUCTS P.O. Box 804 Mahwah, New Jersey 07430 201-934-0022

CIRCLE #101 on Reader Service Card ▲

CONTROL DASD GROWTH

Through Effective Storage Management

By Steve Adil

Corporation data processing budgets have been undergoing increasing scrutiny from financial management in this era of corporate cutbacks. Consequently, budget dollars spent on data storage are important targets for cutbacks as cost control moves place pressure on data processing budgets. Direct-access storage costs can account for a majority of a company's data processing hardware expenditure. As a result, many corporations have begun to consider the idea of developing a storage management structure to control DASD growth.

Controlling DASD growth by developing a storage management structure has merit for several reasons. The most obvious is slowing of DASD growth and corresponding easing of pressure from senior management to reduce costs. Another reason is improving communications between users of data storage and managers of the resource, thus increasing the probability of timely and effective data management activity. It can also provide positioning for evolutionary changes and benefits afforded by the IBM Data Facility Product (DFP) family.

The trick is to plan the purchase of hardware in time for its requirement but not to purchase too much too soon. Buying large quantities of hardware to have on hand for the eventuality of its use is a costly luxury. The cost of funds expended for hardware that will just sit around boosts the real cost of the hardware past a justifiable limit. With Volume Purchase Agreements (VPAs), a customer gets a volume discount based on planned purchases throughout the year and the purchases do not have to come all at once.

How Can Storage Management Control Costs?

The most valuable effect of storage management is the imposition of structure and predictability upon the use of data storage resources. This process allows the storage capacity planner to have a much better idea of the coming requirements for data storage. Planning for data storage capacity is a difficult task due to the complex and highly technical nature of its use. The storage management architecture can provide for the ability to track resource usage without the need for high levels of technical expertise and a large staff of technical experts that can communicate with the various user areas. An effective storage management structure then ensures that the data storage needs of the non-DP areas of the company can be adequately communicated.

The combination of centralization/standardization provided by the application of creative software solutions in the storage management architecture can eliminate the time spent in monitoring resource usage. The proper role of storage management is to eliminate the need for the user to monitor resource usage. A responsive structure can unburden the user from such considerations while effectively and efficiently controlling how the data storage resources are used. The only responsibility the user should have is to accurately communicate his or her needs for data storage to the appropriate area.

In addition to the time wasted there can be a huge expense due to wasted space on the DASD subsystem. The architecture of storage management can eliminate wasted DASD space through a technique usually referred to as "pooling." Pooling as used at Connecticut National Bank was described in detail in the article, "Storage Management's Relationship to DFSMS" (Data-Facility Storage Management Subsystem), in the May/June, 1988 issue of MAINFRAME JOURNAL. It is used to segregate types of data storage based on usage.

Over 1,000
companies

Worldwide

are benefitting from
Empact Software
products including:

**The Software Product
that Stops B37, D37,
E37, Abends.**

**The Write-To-Operator
Message Management System
For MVS Automation.**

**Dynamic Data Set Pooling
and Standards Enforcement.**

Please call for references
and a free evaluation at
(404) 483-8852.

EMPACT software inc.
A Boole & Babbage Company

Corporate Headquarters
1803 Overlake Drive • Conyers, Georgia 30208

CIRCLE #119 on Reader Service Card ▲

The cost-controlling effects of pooling are twofold. Pooling data on several volumes eliminates the ownership of volumes associated with specific groups or projects. The problem with project ownership of volumes is that not all projects require full-volume storage. Because no other project can use the unused space, it remains idle waiting for use by the owner sometime in the future. The multiplication of this unused space across many project-owned volumes can add up to unused amounts of dozens of volumes depending on the size of the shop. DASD pooling groups all of the projects together, eliminating the boundaries of DASD ownership and the resulting wasted space. Pooling can, however, be problematic if not properly controlled. Effective control is established through the techniques in place in the architecture.

FIGURE 1

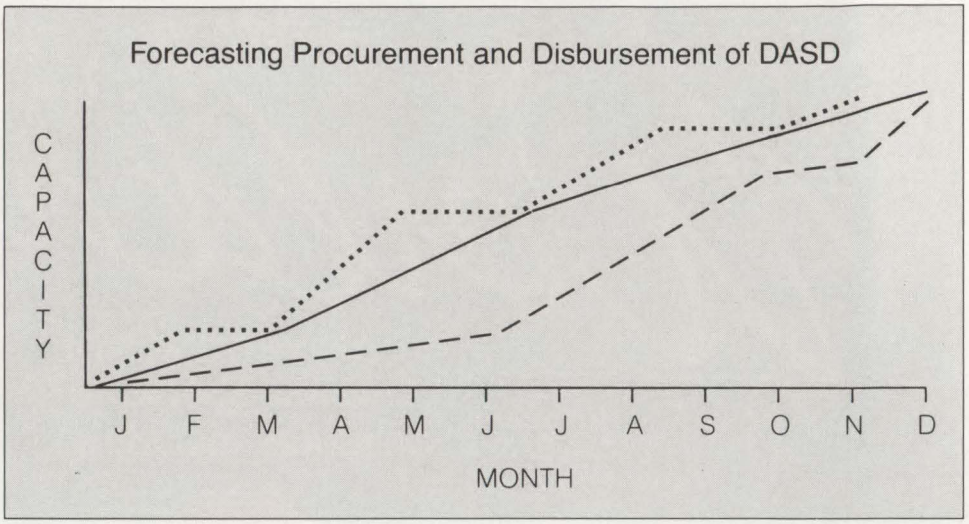

The next consideration of storage management structure is to provide a process for users to communicate DASD requirements and a process to satisfy the requirements. The following is a description of the processes as they exist at Connecticut National Bank. There may be some pieces that you could use in your particular installation.

Forecasting Procurement and Disbursement of DASD

The key to "just in time" procurement lies in encouraging good communication flow between the users of the resource and the providers. It is important to spell out to the users what your expectations are as provider of the resource. It is also important to let the users know what you will provide to them in return. What is needed from users is accurate forecasting of storage needs. What you provide to the users is a guarantee of adequate storage levels. The system that handles this process must also be flexible enough to accommodate the inevitable "unforecasted" requirements for storage.

A measure of how well the provider is doing is the absence of "X-37" (B37, E37 and so on) job abends and timely turnaround of DASD requests. In addition to this, the absence of large amounts of unused DASD space is an indication that the requestors (forecasters) are providing accurate information.

In Figure 1, the dotted line depicts a situation in which storage is purchased later than required. There is likely to be a lot of pain in the user areas as storage frequently becomes tight. The dashed line depicts an overabundance of storage. Users will be happy, but the organization is spending money that it does not have to spend. The straight line is the best way to procure storage.

The process for procuring data storage begins with generation of the annual data processing budget and continues with regular reforecasting exercises.

Budget Generation

During the period August-December, the data is gathered for the coming year's

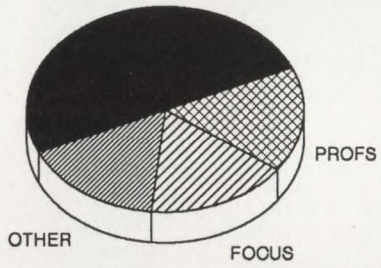

Now that you know that CMS is using most of your VM system...

Don't You Need To Know How Cms Is Being Used?

With CMAP, CMAP/FOCUS and CMAP/PROFS, you can get information on user and command activity, response time analysis and other vital information required for effective performance analysis of your VM system.

For Info Call: 800-527-3237
In California: 619-487-8030

VM/CMS
A subsidiary of
First Alliance Software & Technologies, Inc.

Managing DASD Performance Can Be a Pleasure Trip... Not a Survival Course

DASDMON will revolutionize the way you manage MVS DASD performance.

Traditional DASD tuning techniques are being outpaced by faster CPUs and increasing user demands for more online data. Although cache controllers, data spaces and hiperspaces will relieve some of the strain, they also create new complexities. Tuning DASD can feel like struggling through a survival course.

DASDMON takes you away from it all.

DASDMON offers you system managed performance today. By cutting through the complexities of conventional tuning, managing DASD performance is more automatic and streamlined than ever before. Once you've set response time objectives, just sit back and...

- **DASDMON** will identify problems by continually measuring individual I/O response times and comparing them to your objectives.
- **DASDMON** will determine the cause of problems by gathering and analyzing data. No longer will you spend hours scanning reports.
- **DASDMON** will present solutions to your problems in precise English, and simulate their effects to ensure that a new problem is not introduced.
- **DASDMON** will select the best possible set of solutions so you can implement recommendations with confidence.

And, if you have any questions or concerns along the way, our experienced technical support staff will be glad to guide you—24 hours a day.

For more details, call **(800) 323-2600** or **(412) 323-2600** in Pennsylvania.

DASDMON from Duquesne Systems.
Enjoy the results.

DUQUESNE SYSTEMS
Two Allegheny Center
Pittsburgh, PA 15212

CIRCLE #80 on Reader Service Card ▲

budget. The period covered by the budget is January-December. During the gathering process, any area requiring DASD for the coming year forwards its estimates to the budget planning areas. Any amount of DASD can be requested — from one cylinder to many volumes. In our case, the user thinks in terms of 3380 cylinders or volumes. Whether his data ends up on single, double or triple capacity DASD is irrelevant for planning purposes. The placement of his data will be dealt with when the time comes. Obviously, the budget document will be extremely large and cumbersome if every DASD cylinder is listed therein. Instead, the details of all requests for less than one volume (885 3380-type cylinders) are handled in a separate document and summarized on the budget document as a general "new projects/application growth" category. Projects requiring one or more volumes will appear on the budget document itself along with a separate justification document.

Forecasting

After undergoing management review and approval, the final form of the budget is printed and will be used for planning purposes throughout the year. Each quarter a DASD forecast is sent to the DP division for revision. This is the detail document needed for accurate planning of DASD purchases. DASD will be ordered quarterly to fulfill the storage requirements as outlined in the forecast document. As projects are inevitably delayed or accelerated, so is the purchase of the required DASD. By delaying purchase of DASD until it is actually required, the company can delay expenditures resulting in lower "soft" costs such as the cost of funds. This is in part the theory of "just in time" purchasing methods.

As stated above, the key to a good storage management structure is responsiveness to the users of the resource. By delaying purchase of DASD until it is actually required, the storage management planners run the risk of getting caught without enough storage capability. Several steps can be taken to prevent this from occurring. In a "pooled" storage environment there can be adequate padding in the DASD pools to provide a temporary bridge until new DASD is installed. There are also software solutions available such as STOP X-37 from Empact Software, Inc. (Conyers, GA). STOP X-37 will trap B37, D37 and E37 abends and provide storage for "spilled" datasets or more than sixteen extents, depending on what is required. By providing enough capacity in the STOP X-37 "spill" pool, the increase in X37 abends that occurs as the production DASD pools get tight on available storage will be trapped and no interruption to production processing will occur. This is especially useful because spot shortages can occur even in DASD pools with plenty of available space.

Also key to providing a responsive system is the guarantee of rapid turnaround to budgeted requests for storage. It has been my experience that after a user has been held up waiting for requested storage to be provided, he or she will respond by forecasting more DASD earlier than needed to compensate. Once the user is comfortable that the requests will be satisfied in a timely manner, accurate forecasting is usually provided. Below is a discussion of the disbursement system that has been developed to communicate and fulfill requests in an efficient and timely manner.

Disbursement

At Connecticut National Bank, the vehicle used for processing DASD requests is an electronic mail system. With this particular E-Mail package, electronic "forms" can be created for the three types of DASD requests usually processed. The three types of requests we have found to be the most common are:

- Budgeted DASD requests — a budget number is required to tie the disbursement back to a specific budget item
- Unbudgeted DASD requests — approval is required from application manager, department head and division head before disbursement occurs
- "Loaner" DASD requests — for temporary (30 days or less) requirement of storage. Subject to DASD availability.

The forms are submitted and processed only for requests of one volume or more. This limits the degree of imposition that storage management structure places upon the user. Requests for less than one volume are satisfied even if unbudgeted without being subject to approval. If the request is unbudgeted, it is automatically logged into a report file. Each quarter, storage management provides bank senior management with reports that outline all DASD disbursed for budgeted or unbudgeted projects. The gathering of disbursement data is facilitated by an on-line system that handles the communication of requests from the programming to the data management area. Requests are constructed when programmers fill in ISPF-based screens that ask for background information about the request. The Automated Data Management Requests System (ADMS) translates the information into a data management transaction that informs the data manager of the nature of the request or group of requests (one or more transactions or requests are known as a data management "package"). This software is a good tool for easing the impact of centralizing the data management effort because it actually reduces the effort required by the programmers to maintain the files of a production system. The ADMS software is marketed through Advanced Software Products Group (Naples, FL).

Data management packages are electronically transmitted to the data manager for action and the quality control area for review. After the package is reviewed and accepted by quality control, the data manager can perform the data management

See DASD Growth *page 99*

The Automated Data Management System *ELIMINATES PAPERWORK!*

REDUCE THE AMOUNT OF TIME NECESSARY FOR CONSTRUCTING DATA MANAGEMENT REQUESTS BY NOTIFYING MANAGEMENT THROUGH ADMS TSO/ISPF SCREENS!

CREATE - TRANSMIT & PRINT REQUESTS WITH ADMS --- EVEN VIEW HISTORICAL REQUEST DATA!

IT'S AS EASY AS A-D-M-S

ASPG

2335 Tamiami Trail North
Suite 401
Naples, Florida 33940

1-800-662-6090

If...

One product and one hour could double VSAM performance

VSE UNDER THE COVERS

A Look at the Internals of the VSE Operating System

By Eric L. Vaughan

As IBM's leading operating system in terms of number of installations, VSE enjoys a wide variety of uses and applications.

The internals of the VSE operating system are mapped out into major control blocks with reliable anchor pointers. The relative ease of debugging and implementation of user modifications are important contributors to the longevity and popularity of the VSE operating system.

VSE does what every good operating system should do; manage I/O and CPU resources. Dozens of control blocks live "under the covers" of VSE. These control blocks are like old friends during the process of problem determination. Throughout many years and several releases of DOS/VSE, the control blocks have changed and expanded. For purposes of our peek inside, this article will reference the control blocks relative to VSE/AF 2.1.

Before problem determination and debugging can begin, a full understanding of these control blocks is required. We will examine the control blocks in terms of four major functional areas: system anchors, partition related, task related and I/O control.

Trusted Friends: the System Anchors

Certainly one of the first anchors to examine in the course of problem determi-

FIGURE 1

SYSTEM COMMUNICATIONS REGION (SYSCOM)
KEY AREAS

HEX DISPLACEMENT	DESCRIPTION
00-03	Address of Error Block
04-07	Hard Wait Code
0C-0F	Pointer to SYSRES PUB
1C-1F	Address of LTA
20-23	Beginning of Problem Program Storage
24-27	Address of CHANQ
2C-2D	Number of Partitions
30-33	Total Virtual Storage Size
3C-3F	Address of VIO Communication Area
58-59	Task ID of LTA Owner
5A-5B	Task ID of Last Task Dispatched
7C-7F	Address of Job Accounting Common Area
9C-9D	VTAM's PIK
9E-9F	POWER's PIK
A0-A1	ICCF's PIK
DC-DF	Address of SMCB Address Table
EC-EF	End of Real Storage (370 mode only)
F5-F7	Address of SVA
F8-FB	Address of System Getvis Area
130	End of SYSCOM

SYSCOM is pointed to at fixed location x'80'.

Then...

Introducing HYPER-BUF, the VSAM buffer allocation expert

- **Shatters VSAM performance barriers**
  — Customers consistently report 30% to 80% reductions in batch run times
- **Optimizes CICS buffers**
- **Installs in less than one hour**

Other VSAM products from Goal Systems, the Original VSAM Experts

/* **FAVER —** A VSAM file management tool with comprehensive reorganization and backup/restore capabilities.
VSAMAID — Tunes VSAM datasets automatically while eliminating inefficiencies. A cluster modelling, performance, and capacity planning tool.
MASTERCAT — An online and batch VSAM catalog navigation tool.
COMPRESSOR/MVS — Provides compression of VSAM datasets. Can triple disk capacity. Available for MVS.

To Find Quality Call 800-848-4640.

All products, except COMPRESSOR/MVS, are available for VSE and MVS.*/

Goal Systems

Goal Systems International • 7965 N. High Street • Columbus, Ohio 43235 • 800-848-4640
Goal Systems International S.A.R.L. • 88 avenue de Wagram • 75017 Paris, France • Phone: (1)42 67 55 55 • Telex: 641.094

CIRCLE #79 on Reader Service Card ▲

nation is the System Communications Region or SYSCOM. The address of SYSCOM can always be located at fixed location x'80'. SYSCOM provides many pointers to important tables as well as critical configuration and status information (see Figure 1). Information such as the number of partitions in the system, the last task set active by the supervisor and the system hard wait code are examples of the data available here. During debugging of a VSE system hard wait for example, the last task dispatched at +x'5A' into SYSCOM should be examined to gain insight into the events leading to the system's sudden demise.

The second important system area is the Fixed Control Area or PSA (see Figure 2). The PSA, anchored at fixed storage location x'240', contains several vital pointers relating to the current task in control such as the current TCB, TIB, PIB and PCB. Additionally, the Partition Selection String (PSS) and the address of the PCB address table arranged in priority order (PPRTYOWN), both critical instruments used in the science of task selection, are found here. The PSA is an abundant source of information during problem determination.

An additional, yet partition-related anchor can be found at fixed location x'14'. The Communication Region (COMREG) of the current partition is anchored here. The COMREG contains addresses and pieces of information relative to the active partition.

FIGURE 2

FIXED CONTROL AREA (PSA)

LOW CORE ADDRESS (HEX)	DESCRIPTION
240-241	Flags
242-243	Routine Identifier (RID)
244-247	Current Time Pointer
248-24B	Address of PIB2TAB
24C-24F	Address of System PCB
250-253	Address of SCB Address Table
254-257	Address of Current SCB
258-25F	Partition Selection String (PSS)
260-263	Pointer to Current TCB
264-267	Pointer to Current TIB
268-26B	Pointer to Current PIB
26C-26F	Pointer to Current PCB
270-273	Debug Header Address
274-277	Debug Parameter Address
278-27B	Address of First Level Interrupt Handler
27C-27F	Address of Dispatcher
280-2BF	ERA Save Area
2C0-2C3	Address of TIBATAB-4
2C4-2C7	Address of PCB Address Table
2C8-2C9	Number of Page Frames Available
2CA-2CB	Minimum Number of Pages Not Fixed
2CC-2CD	Save Field for RID

FIGURE 3

PARTITION IDENTIFICATION KEY (PIK) ASSIGNMENTS

NPART	BG	FB	FA	F9	F8	F7	F6	F5	F4	F3	F2	F1
12	10	20	30	40	50	60	70	80	90	A0	B0	C0
11	10		20	30	40	50	60	70	80	90	A0	B0
10	10			20	30	40	50	60	70	80	90	A0
09	10				20	30	40	50	60	70	80	90
08	10					20	30	40	50	60	70	80
07	10						20	30	40	50	60	70
06	10							20	30	40	50	60
05	10								20	30	40	50
04	10									20	30	40
03	10										20	30
02	10											20

The PIK of a partition is located at x'2E-2F' of its COMREG.

System anchors to examine are SYSCOM, PSA AND COMREG.

Each One Has Four

Within the VSE system there are four major control blocks associated with each partition. The Program Information Block (PIB), the Program Information Block Extension (PIB2), the Partition Control Block (PCB) and the COMREG. The secret to locating a particular partition's control blocks lies in the usage of the Partition Identification Key or PIK. The PIK always remains the same for the background partition. It varies, however, for all foreground partitions based on the number of partitions generated in the system (see Figure 3).

A partition's COMREG can be addressed by its PIB2 + x'00'. The COMREG contains additional system pointers and partition information (see Figure 4) such as the current date, the active job name and most importantly, the PIK of the partition.

The PIB and PIB2 tables are anchored at x'5A' and x'7C' into COMREG respectively. Additionally, the PIB2 table is pointed to by a full word at low core x'248' within the PSA.

Computer Associates introduces total data center automation for VSE.

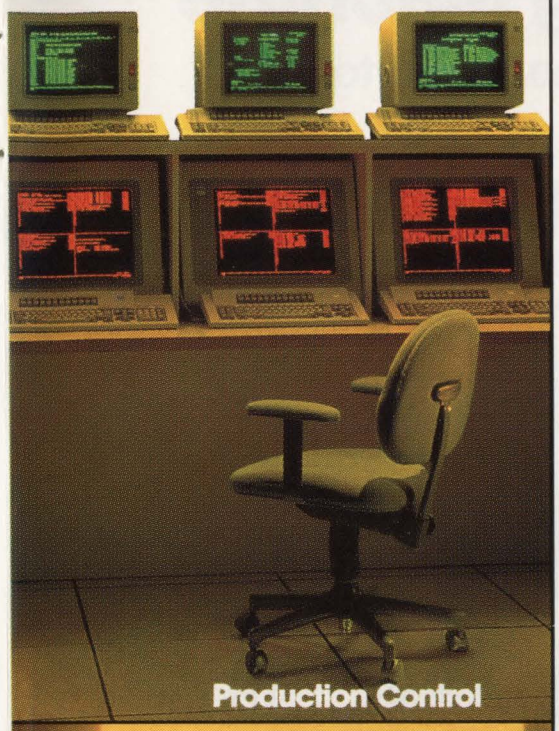

Computer Associates brings you the industry's only complete line of VSE systems software: **CA-UNICENTER®/II VSE**, the total solution for data center automation.

With **CA-UNICENTER/II VSE**, you gain extensive tape and disk management capabilities, automated production scheduling and output distribution, performance measurement and job accounting facilities, the industry's leading security software, essential sorting and report writing utilities and even more.

Through advanced integration and automation, **CA-UNICENTER/II VSE** ensures efficient utilization of data center resources, virtually eliminates manual involvement and dramatically improves overall productivity of the data center—enabling you to achieve levels of efficiency never before possible.

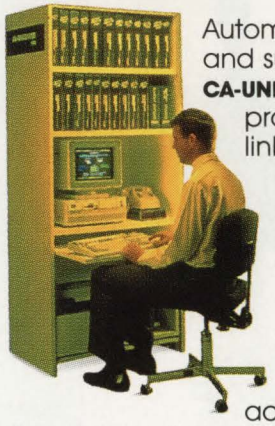

Automated service and support, too: **CA-UNISERVICE®/II** provides a secure link between your mainframe and Computer Associates' Customer Service System, 24 hours a day. You get online access to software fixes, interactive problem resolution, product tutorials and more.

Call Dana Williams today:
800-645-3003 (Ext. 5811).

© 1988 Computer Associates International, Inc.
711 Stewart Ave., Garden City, N.Y. 11530-4787

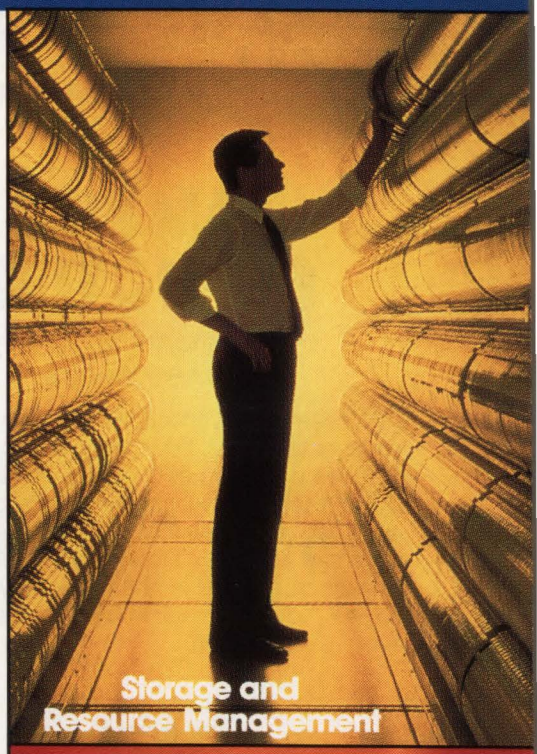

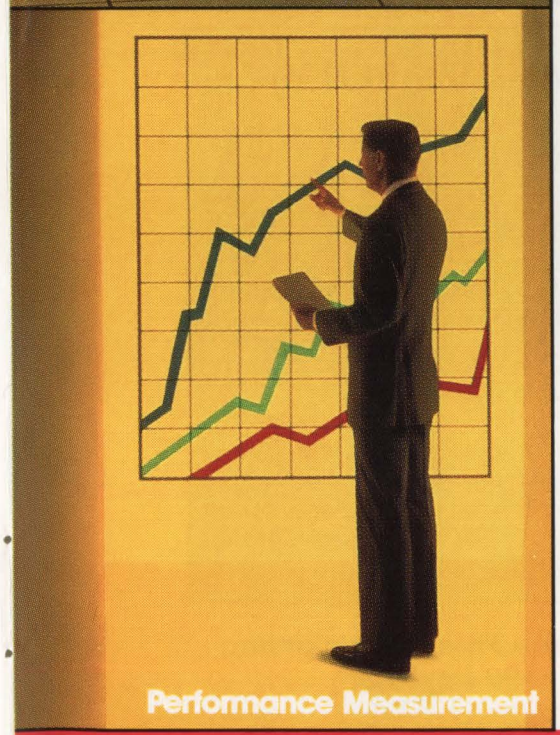

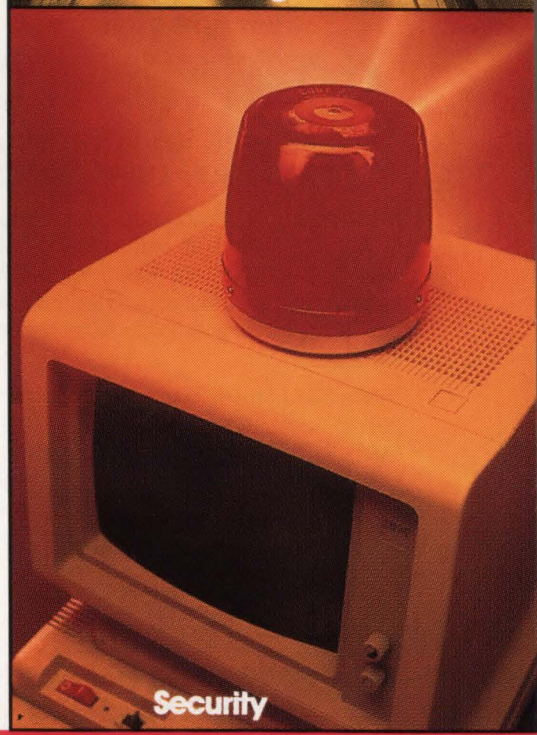

COMPUTER® ASSOCIATES
Software superior by design.

- World's leading independent software company.
- Broad range of integrated business and data processing software for mainframe, mid-range and micro computers.
- Worldwide service and support network of more than 100 offices.

Resource & Operations Management • Financial • Banking • Graphics • Spreadsheets • Project Management

CIRCLE #105 on Reader Service Card ▲

IAM vs VSAM

The Incredible Shrinking Machine

IAM Reduces the Size of Your VSAM Files by 30 to 70%

IAM FILE STRUCTURE

IAM SAVES 20 to 40% DASD SPACE

IAM uses an advanced file structure which is far superior to VSAM. IAM's supercompressed index requires a fraction of the space taken by VSAM. IAM's freespace concepts make much more efficient use of disk space. IAM's blocksizes are not restricted as VSAM's are, making full utilization of each track. IAM is not affected by large key sizes which can result in VSAM wasting CI's in every Control Area.

DATA COMPRESSION

SAVES AN ADDITIONAL 20 to 50% DASD SPACE

IAM optionally compresses data records. Most files contain records with unused fields or repeating sets of characters. When IAM applies its proprietary compression techniques, the result is an additional 20 to 50% reduction in file size.

IAM's CPU time is dramatically less than competing compression products. In fact, since IAM's CPU time is normally much less than VSAM, IAM with data compression takes less CPU time than normal VSAM processing.

RELEASE

AUTOMATIC RELEASE OF UNUSED SPACE

IAM takes the guessing game out of VSAM space allocation. Large amounts of disk space are wasted when users over-estimate how much space VSAM requires or how many records a file will contain. VSAM cannot release overallocated space.

TRANSPARENT

VSAM files account for the lion's share of disk space used in most installations. Online systems (CICS), BATCH jobs, TSO, SMP/E and other applications make extensive use of keyed index VSAM (KSDS) files.

IAM is a transparent alternative to VSAM KSDS files, which substantially reduces the impact of VSAM processing in your installation. There are no modifications to programs or JCL to use IAM files in place of VSAM.

SMF ANALYSIS

VSAM SIZE REPORT

DATA SET NAME	ALLOC TRKS	TOTAL EXCPS
BIG.CLUSTER	37155	2507803
BIG.CLUSTER.DATA	37100	2105001
BIG.CLUSTER.INDEX	55	402802
A.FILE.SMALLER	16540	679216
A.FILE.SMALLER.DATA	16500	270501
A.FILE.SMALLER.INDEX	40	408715
SMPE.TDFP223.CSI	12315	3880211
SMPE.TDFP223.DATA	12300	3075021
SMPE.TDFP223.INDEX	15	805190

IAM's SMF analysis program identifies your largest and most highly used VSAM files. **To see your VSAM usage, send for the FREE SMF reporting program (IAMSMFVS).**

Call for a Free No Obligation 90 Day Trial

From the Makers of FDR & ABR
Supports MVS and MVS/XA

INNOVATION DATA PROCESSING

Innovation Plaza, 275 Paterson Avenue
Little Falls, NJ 07424 (201) 890-7300

CIRCLE #97 on Reader Service Card ▲

Both the PIB and PIB2 entries are 16-byte entries. This is the scheme behind each PIK as a multiple of 16. Given the PIK and the anchor pointer, you can easily index to a partition's control blocks. Historically, both the PIB and PIB2 are old tables maintained for sake of compatibility. The PIB2 has a few important pointers, however, such as the address of COMREG at +x'00', the Task ID (TID) of the main task at +x'04' and the PCB pointer at +x'08'.

The PCB is a critical control block used during task selection by the supervisor's dispatcher (see Figure 5). This control block was significantly reworked for VSE/AF 2.1 to include such information as the current Space Control Block (SCB) pointer in a Virtual Addressability Extension (VAE) system. The PCB also contains the pointer to the main task save area where PSW and registers are stored upon an interrupt. The Task Selection String (TSS) and Task ID String (TIDSTR), both essential components used during task selection, are located within this control block. During problem determination, these areas will aid in determining whether a task is ready to run or is waiting on an event or resource.

VSE's control logic breaks the "work unit" down further from partitions to tasks. One or more tasks may be active in a partition at any one time. All resource scheduling is based at the task level within the highest priority partition.

A Tisket, A TASKet?

There are two control blocks for each main and sub-task in the VSE system; the Task Information Block (TIB) and the Task Control Block (TCB). These control blocks are no longer necessarily found in the supervisor's storage, rather they are allocated dynamically at task initiation and are used to control and monitor each task's status. The unique TID of each task is used to locate the blocks.

VSE's control block architecture has grown to accommodate new facilities.

Remember the TID of the last task dispatched found at SYSCOM +x'5A'? Often times this is where we start to work into the control blocks. All of the TIDs within a partition are located at TIDSTR in the PCB.

TIBs can be located by using a table called the TIB Address Table (TIBATAB). The TIBATAB is a table of four-byte address pointers to the TIB for each task in the system in TID sequence. Using the TIBATAB pointer located at x'2C0' in the PSA (actually this is the address of TIBATAB minus four to allow for easy indexing), the TIB for a given TID can

FIGURE 4

COMMUNICATIONS REGION (COMREG)
KEY AREAS

HEX DISPLACEMENT	DESCRIPTION
00-07	Job Date
17	UPSI Byte
18-1F	Job Name
20-23	Problem Program End Address
2E-2F	PIK of Partition
40-41	Address of PUB Table
4C-4D	Address of LUB Table
5A-5B	Address of PIB Table
6C-6D	Address Space ID (370 mode only)
70-73	Address of Sysparm
74-77	Address of Job Accounting Table
8F-96	Procedure Name
A0-A3	Pointer to POWER PCB
E4	End of COMREG

The Current COMREG can be found at low core address x'14'.

FIGURE 5

PARTITION CONTROL BLOCK (PCB)
KEY AREAS

HEX DISPLACEMENT	DESCRIPTION
00-01	Length of this PCB
04-07	Partition Priority Mask for PSS
14-15	PIK
16-17	Length of the TSS (For TRT)
18-1B	Task Selection Status (TSS)
20-43	Task Id in Priority Sequence
4C-4F	SCB Pointer of Active Space
54-57	Address of Main Task Save Area
5C-5F	Virtual Partition Begin Address
60-63	Virtual Partition End Address + 1
64-67	Real Partition Begin Address
68-6B	Real Partition End Address + 1

The PCB can be located at PIB2 + x'08'.

FIGURE 6

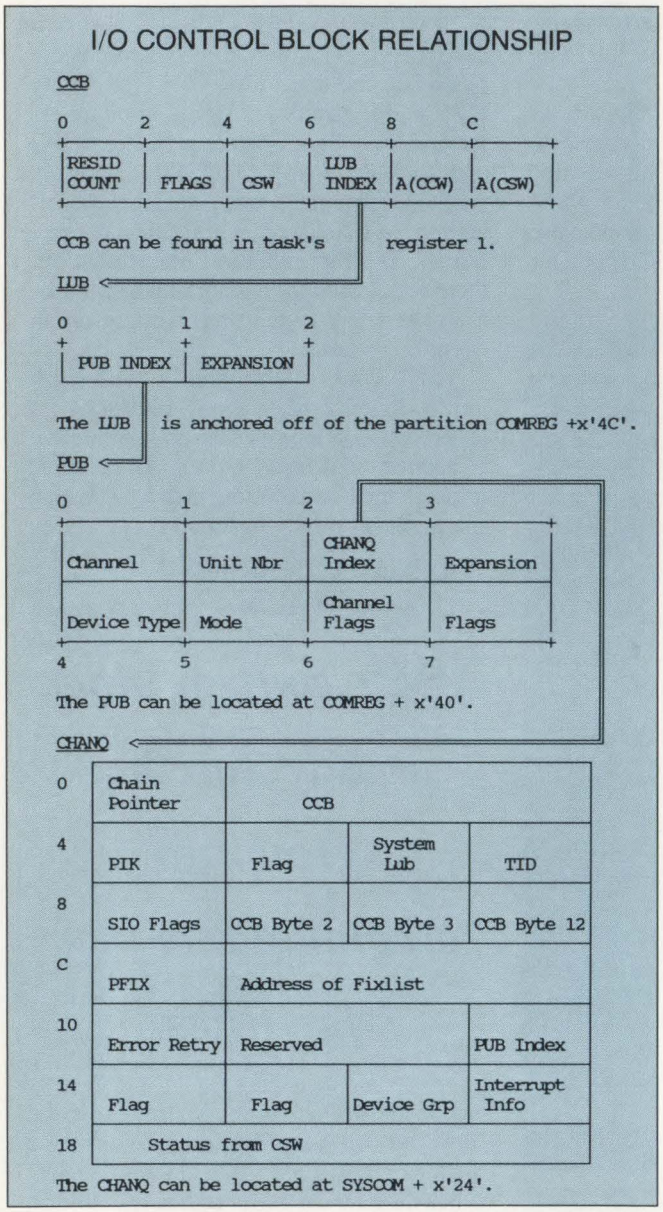

FIGURE 7

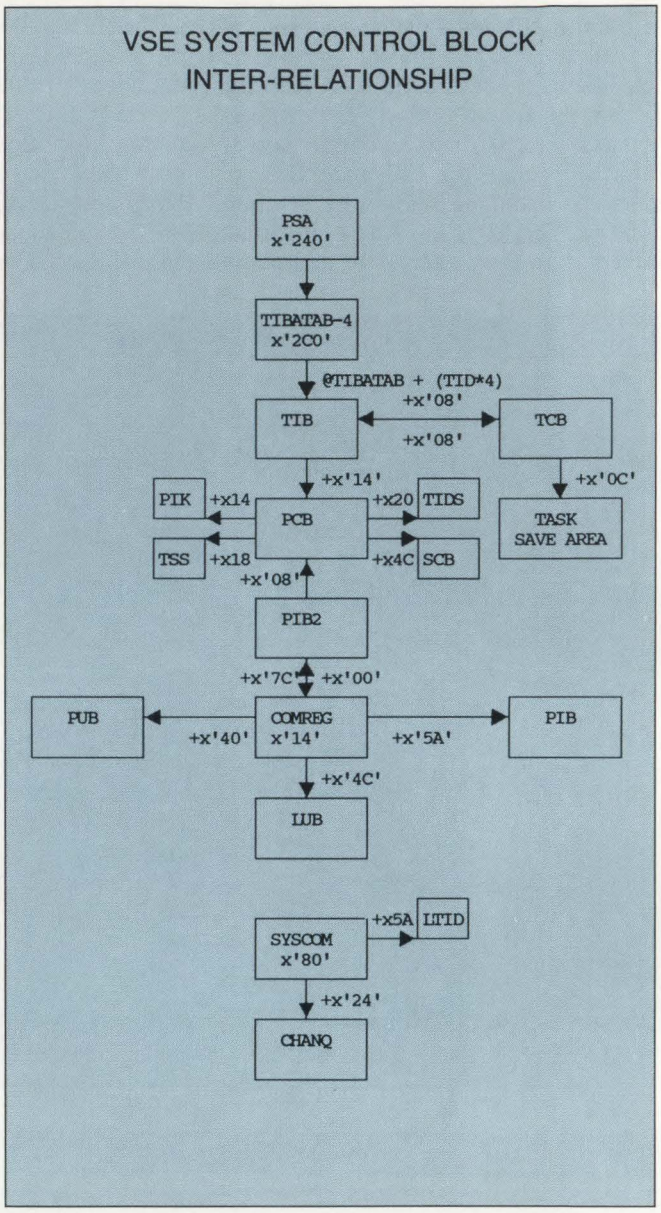

be calculated easily. The TIB is heavily used by the dispatcher during task selection and wait processing. It points to the TCB for the task at displacement x'08'. A cancel code for a task is posted in the TIB and all information required to analyze a task wait can be found here.

The task's TCB can be variable in length depending on several characteristics such as whether the task is a main task, sub-task or system task. The TCB anchors the TIB, likewise at x'08'. This control block is vital during debugging as the task's save area is pointed to at x'0C'. The TCB is also used by the Fetch routine and for CCW translation during I/O processing.

What Goes In, Must Come Out

The management of I/O resource scheduling is handled through the use of control blocks as well.

There are four major I/O related control blocks: the Command Control Block (CCB), the Logical Unit Block (LUB), the Physical Unit Block (PUB) and the Channel Queue (CHANQ).

The CCB is a control block originating within a task's storage area. The CCB represents a request for I/O and contains information such as the pointer to the particular I/O instruction (that is, read/write) to be executed and to the target device. For debugging purposes, if a task is waiting on I/O completion, the task's general register one should be examined to locate the CCB. Optionally, the CCB will be pointed to in the task's TIB +x'04' or TIBSTATE.

There is a LUB table for each partition consisting of a two-byte entry for each logical unit that can be assigned by a problem program. This table is used by the supervisor to calculate the target I/O device as indicated in the requestor's CCB. The LUB table can be located at x'4C' into the partition COMREG. The LUB table consists of a one-byte index (hence the arbitrary 254-device limit) to point to the actual device entry in the PUB table.

Lee Veal On VSE Console Support...

"Anything IBM gives you, DOCS makes five to ten-fold better."

DOCS™

With 17 years systems software experience, Lee Veal, Software Support Supervisor for The City Of Garland, Texas, has seen a lot of software come and go. But there's one product Lee won't let go—DOCS (Display Operator Console Support) from SMARTECH Systems. And here's why...

I recommended DOCS because it looked like it was going to save us a lot of money. With the ability to call up the VSE system console from multiple terminals, we're able to de-centralize our access. And that cuts down on the phone calls received by the operational staff, because programmers have access to DOCS in their building ...and at home.

How easy it is to install DOCS...

It's not like re-inventing the wheel to bring DOCS up. You don't have to change your operational layout or procedures just because DOCS is in place. In fact, it actually makes the system come up faster.

Plus DOCS allows the terminals to be shared back and forth between CICS or any other on-line software. They can be dual purpose—and we need those terminals for other functions.

On DOCS security features...

DOCS allows us to use password securities and "read only" type securities whenever the DOCS consoles come up. We can even include or exclude partitions on each individual console. We use these features in the programmer area—for the terminals that the programmers have access to.

How DOCS Dynamic System Status Display (DSSD) increases efficiency...

It tells us if a job is running or why it stopped. This makes a difference when we have a problem. We can know what the program is waiting on and address the problem saving us a lot of time. If we didn't have this, we'd be "flying blind."

How DOCS reduces keying time...

With DOCS we're able to set up our PF-Keys on-line to store commonly used responses like IGNORE, CANCEL and RETRY. Also, there are features that let us "pre-answer" questions. One we use more than any is the "call back" feature where you can call back your previous reply and answer that same thing again for the next response. Plus, we have the ability to recall any line on the screen for input. It saves the operators time from sitting and pecking out those characters.

On the need for DOCS in today's environment...

It simply makes your operation run easier with it than without it. Most of the benefits of DOCS you really can't say in words how good they are—you just have to experience it.

Why not experience it for yourself? Try DOCS FREE for 30 days. It takes just 30 minutes to install. Call us for more information.

1-800-53-SMART

Outside the U.S., call (214) 630-8324

SMARTECH SYSTEMS, INC.

Turning high technology into *SMART TECH*nology.℠

10015 W. Technology Blvd., Dallas, TX 75220
Telex: 9102503110

VSE and IBM are registered trademarks of the International Business Machine Corp. SMARTECH and DOCS are trademarks of SMARTECH Systems Inc. Copyright © 1988 SMARTECH Systems, Inc. All rights reserved.

CIRCLE #65 on Reader Service Card ▲

The PUB table consists of one eight-byte entry for each device and is anchored at x'40' into the COMREG. The information in this entry includes the device address, type and status of any I/O outstanding. Additionally, a one-byte index into the CHANQ table can be found if any outstanding I/O exists for the device.

The CHANQ, anchored off of SYSCOM +x'24', contains a 32-byte entry for each I/O request. The information stored here includes the CCB (relocated if necessary), TID and PIK of the original requestor.

Figure 6 illustrates the correlation of all of the I/O related control blocks.

Put It All Together

In Figure 7, the overall structure and relationship these control blocks have with each other is illustrated. The essence of using control blocks is to provide re-entrancy and ease of resource management. Common routines have been written to handle a single resource function. By using different control blocks, a variable amount of resources can be managed and expansion is an inherent benefit as there are no explicit limitations.

The control block architecture of VSE is solid and has grown over the years to accommodate new facilities. In the next article, the dispatching and scheduling techniques employed by VSE and the interactions with these control blocks will be examined.

ABOUT THE AUTHOR

Eric Vaughan is president of Smartech Systems, Inc., a systems software firm specializing in the VSE market. Vaughan has more than nine years of technical experience with VSE and VM operating systems and has worked extensively in the field of telecommunications. Smartech Systems, Inc., 8701 Carpenter Freeway, Dallas, TX 75247, (214) 630-8324.

U.S. Postal Service
STATEMENT OF OWNERSHIP, MANAGEMENT AND CIRCULATION
(Required by 39 U.S.C. 3685)

1A. TITLE OF PUBLICATION MAINFRAME JOURNAL	1B. PUBLICATION NO. 0 8 9 5 5 7 5 1	2. DATE OF FILING OCTOBER 27, 1988
3. FREQUENCY OF ISSUE Bi-monthly	3A. NO. OF ISSUES PUBLISHED ANNUALLY 6	3B. ANNUAL SUBSCRIPTION PRICE 0 (Controlled)

4. COMPLETE MAILING ADDRESS OF KNOWN OFFICE OF PUBLICATION (Street, City, County, State and ZIP+4 Code) (Not printers)
10935 Estate Lane-Suite 375 Dallas, Texas 75238

5. COMPLETE MAILING ADDRESS OF THE HEADQUARTERS OF GENERAL BUSINESS OFFICES OF THE PUBLISHER (Not printer)
Same as #4

6. FULL NAMES AND COMPLETE MAILING ADDRESS OF PUBLISHER, EDITOR, AND MANAGING EDITOR (This item MUST NOT be blank)

PUBLISHER (Name and Complete Mailing Address)
Robert H. Thomas 10935 Estate Lane-Suite 375 Dallas, Texas 75238

EDITOR (Name and Complete Mailing Address)
Carol M. Hoag 10935 Estate Lane-Suite 375 Dallas, Texas 75238

MANAGING EDITOR (Name and Complete Mailing Address)
N/A

7. OWNER (If owned by a corporation, its name and address must be stated and also immediately thereunder the names and addresses of stockholders owning or holding 1 percent or more of total amount of stock. If not owned by a corporation, the names and addresses of the individual owners must be given. If owned by a partnership or other unincorporated firm, its name and address, as well as that of each individual must be given. If the publication is published by a nonprofit organization, its name and address must be stated.) (Item must be completed.)

FULL NAME	COMPLETE MAILING ADDRESS
Thomas Publications, Inc (A Texas Corp.)	10935 Estate Lane-Suite 375 Dallas, Texas 75238
Robert H. Thomas (50% of Stock)	10031 Ferndale Rd. Dallas, Texas 75238
Martha W. Thomas (50% of Stock)	10031 Ferndale Rd. Dallas, Texas 75238

8. KNOWN BONDHOLDERS, MORTGAGEES, AND OTHER SECURITY HOLDERS OWNING OR HOLDING 1 PERCENT OR MORE OF TOTAL AMOUNT OF BONDS, MORTGAGES OR OTHER SECURITIES (If there are none, so state)

FULL NAME	COMPLETE MAILING ADDRESS

9. FOR COMPLETION BY NONPROFIT ORGANIZATIONS AUTHORIZED TO MAIL AT SPECIAL RATES (Section 423.12 DMM only)
The purpose, function, and nonprofit status of this organization and the exempt status for Federal income tax purposes (Check one)
(1) ☐ HAS NOT CHANGED DURING PRECEDING 12 MONTHS ☐ (2) HAS CHANGED DURING PRECEDING 12 MONTHS (If changed, publisher must submit explanation of change with this statement.)

10. EXTENT AND NATURE OF CIRCULATION (See instructions on reverse side)	AVERAGE NO. COPIES EACH ISSUE DURING PRECEDING 12 MONTHS	ACTUAL NO. COPIES OF SINGLE ISSUE PUBLISHED NEAREST TO FILING DATE
A. TOTAL NO. COPIES (Net Press Run)	53,934	57,800
B. PAID AND/OR REQUESTED CIRCULATION 1. Sales through dealers and carriers, street vendors and counter sales	-0-	-0-
2. Mail Subscription (Paid and/or requested)	40,150	50,345
C. TOTAL PAID AND/OR REQUESTED CIRCULATION (Sum of 10B1 and 10B2)	40,150	50,345
D. FREE DISTRIBUTION BY MAIL, CARRIER OR OTHER MEANS SAMPLES, COMPLIMENTARY, AND OTHER FREE COPIES	13,434	7,180
E. TOTAL DISTRIBUTION (Sum of C and D)	53,584	57,525
F. COPIES NOT DISTRIBUTED 1. Office use, left over, unaccounted, spoiled after printing	350	275
2. Return from News Agents	-0-	-0-
G. TOTAL (Sum of E, F1 and 2—should equal net press run shown in A)	53,934	57,800

11. I certify that the statements made by me above are correct and complete — SIGNATURE AND TITLE OF EDITOR, PUBLISHER, BUSINESS MANAGER, OR OWNER: Robert H. Thomas, Publisher/Owner

PS Form 3526, Dec. 1985 (See instruction on reverse)

DON'T WASTE ANY MORE TIME WAITING FOR REPORTS

Access POWER and VM queues, and now On-Line Spooling with DTA/PRINT

Avoid the delays of manual report distribution. Any report in any POWER or VM queue can be routed to any CICS attached printer or viewed on any CRT. Security features help you control access to reports.

Your CICS applications can generate reports that can be automatically or manually routed by DTA/PRINT to any CICS printer, or sent to POWER LST, RDR, or PUN Queues.

The headache of Broadcasting reports is ended. Entire reports or parts of reports can be distributed to several locations with a single print request.

Routing reports as they are generated is easy. DTA/PRINT can automatically route reports whenever they appear in user selected queues. Routing can be established by 12 different selection criteria.

Over 100 on-line HELP screens make DTA/PRINT easy to learn and easy to use.

For details and a free 30-day evaluation, call toll free 800-521-6773 or (in Minnesota) 612-591-6100

DTA® Davis, Thomas & Associates, Inc.
Software Products Group
550 Waterford Park
505 N. County Road 18
Minneapolis, MN 55441
800-521-6773 or 612-591-6100

CIRCLE #129 on Reader Service Card ▲

PROBLEMS:

The DOS/VSE Label Area is a performance bottleneck.
Slow disk, relative to CPU, limits performance.

SOLUTION: BIM-VIO

The DOS/VSE 'Virtual' Disk Drive
and Resident Label Area Product.

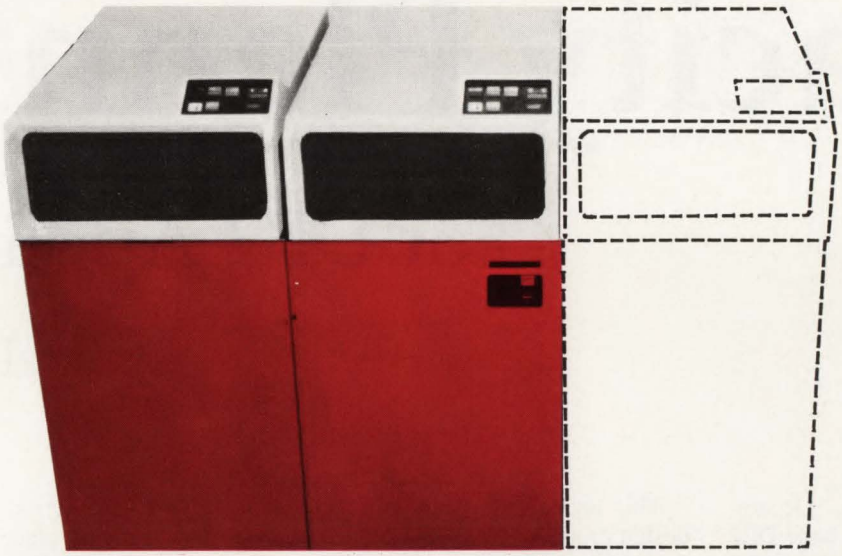

Think of it as a disk drive with zero seek time and rotational delay, having no electrical, air conditioning, or floor space requirements!

BIM-VIO uses a facility in DOS/VSE/SP called "VIO" to map all I/O requests for a non-existent disk address to a virtual memory area outside the normal VSE address space areas. Since this area is in virtual storage, references to it are satisfied at CPU speeds and no actual disk I/O takes place, except possibly if memory paging results. The net result is a potentially significant performance improvement of programs using this area.

The virtual disk can be used for almost anything that does not require permanent retention beyond system start-up (IPL). Examples are compiler work files, SORT work files, temporary files used within or between application jobs. Application programs are not affected, the JCL is simply changed to point to the virtual disk drive 'address'.

A built-in aspect of the product is that the DOS/VSE Label Area is relocated to the virtual disk. This area is one of the most frequently accessed in most DOS sites, so moving it to the virtual disk should result in significant performance improvement to the overall system, regardless of any other specific use of the virtual disk capability.

BIM-VIO is priced at $3600 for a permanent license, $1800 on an annual lease or $180 on a monthly rental.

B I Moyle Associates, Inc. has been dedicated to providing cost effective software solutions, which improve system performance and user productivity, for over 10 years. For more information on BIM-VIO or any of our other quality software products and services, call Jim Kingsbury at 612-933-2885 today.

BIM

B I MOYLE ASSOCIATES, INC.
5788 Lincoln Drive
Minneapolis, MN 55436

612-933-2885
Telex 297 893 (BIM UR)
Member Independent Computer Consultants Assn.

CIRCLE #66 on Reader Service Card ▲

Design, Performance & Capacity Planning

For DB2 Applications

By Joel Goldstein

The time is 10:30 A.M., December 7, 1988. Your first large DB2 application has gone into production. The CPU utilization of your 3090-400E processor has increased from its previous average of 40 percent busy, to more than 90 percent. The response time for all the other on-line systems has quadrupled, the users are screaming and a vice president wants someone's scalp.

Does this sound far-fetched or does it awaken memories of early implementations of other on-line database systems? Many of us have suffered through such a disaster in the past (hopefully distant) and have learned how to plan for new applications and to model their processing requirements. Most of the non-relational DBMS' have been around for quite a while and we know how they perform and how they process their databases. We know how to monitor and estimate the I/Os and resource requirements for each on-line program.

DB2 is a new environment and some of our existing methodologies will have to be modified. Additionally we are not able to control access paths as with older DBMS, nor are we able to monitor *real* I/Os from a query since DB2 uses the VSAM Media Manager directly.

This article addresses the application design process, the tools used to evaluate the performance of specific SQL requests against prototype tables and software tools that may be used to model and project the system resource requirements of DB2 applications.

The data presented here is from DB2 Release 1.3. The next version of DB2 2.1 is scheduled for general availability in October. My previous article in the September/October, 1988 MAINFRAME JOURNAL addressed some of the functional and performance improvements that are expected in the new version. Naturally, the performance improvements will have to be measured as well as the impact of the new features.

Application Design Process

The design process for DB2-based applications is similar to that of any other on-line application using database management software. The critical items are: the data and its volume, the organization/structure of the data, the on-line and batch processing requirements and the expected on-line transaction volume.

Batch processing requirements are a significant part of the overall design process, especially if the system is expected to be available most or all of a twenty-four hour period. Extended operation requirements place difficult burdens upon some otherwise normal activities:

- Table loads of large numbers of daily, weekly, or monthly data records
- Large maintenance runs such as deleting/inserting/updating more than 10 percent of a table (the significance depends upon the size of the table)
- Large reporting requirements that must scan entire tables and will impact on-line terminal response times
- Tablespace/Table reorganizations that are necessary in high update/delete systems
- DB2 systems may require frequent use of the RUNSTATS utility and rebinding of PLANS to enable the OPTIMIZER to select efficient access.

One of the most critical steps in the application development process is the design of the databases (DB2 tables). The theory of database design, especially for relational systems, dictates that the data should be reduced to fifth normal form prior to starting the actual database and table design. The foundations and reasons (elimination of data redundancies) for this process have been set forth by E. F. Codd, C. J. Date and James Martin. The realities of the present "state-of-the-art" of software and hardware will have to bear

You don't have to wait on tables.

Meet tableBASE™. It will improve programmer productivity and system performance — **often by more than 90%.**
A powerful tool for automating table handling in mainframe applications, tableBASE also effectively manages data in memory, dramatically reducing run-time and I/O.
Stop waiting on applications . . . these customers did!

East Midlands Electricity Board (UK)
Cash Management Application, 420,000 transactions

	VSAM	tableBASE	Savings
CPU Time	6.5 min.	3.5 min.	46%
Elapsed Time	2.5 hrs.	5 min.	95%

Security Pacific National Bank
Electronic Funds Transfer system, 660,000 transactions

	VSAM	tableBASE	Savings
CPU Time	14 min.	8 min.	43%
Elapsed Time	7.5 hrs.	28 min.	93%
EXCP'S	895,000	760	99%

Blue Cross and Blue Shield of Ohio
Claims Processing Application, 175,000 transactions

	VSAM	tableBASE	Savings
CPU Time	8 min.	10 sec.	98%
Elapsed Time	30 min.	5 min.	83%

Find out how tableBASE can help you.

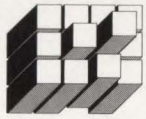

tableBASE

More benchmarks and additional information are available by contacting:
Data Kinetics, 97 Norman Street, Ottawa, Canada K1S 3K5
Tel: (613) 238-6709 FAX: (613) 238-2852

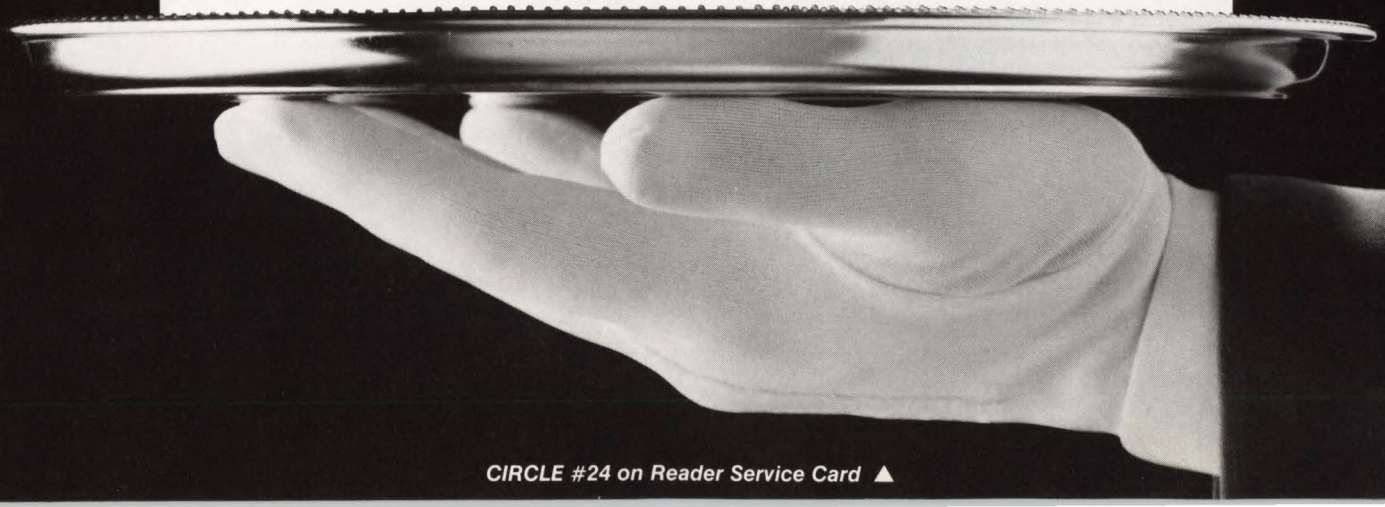

CIRCLE #24 on Reader Service Card ▲

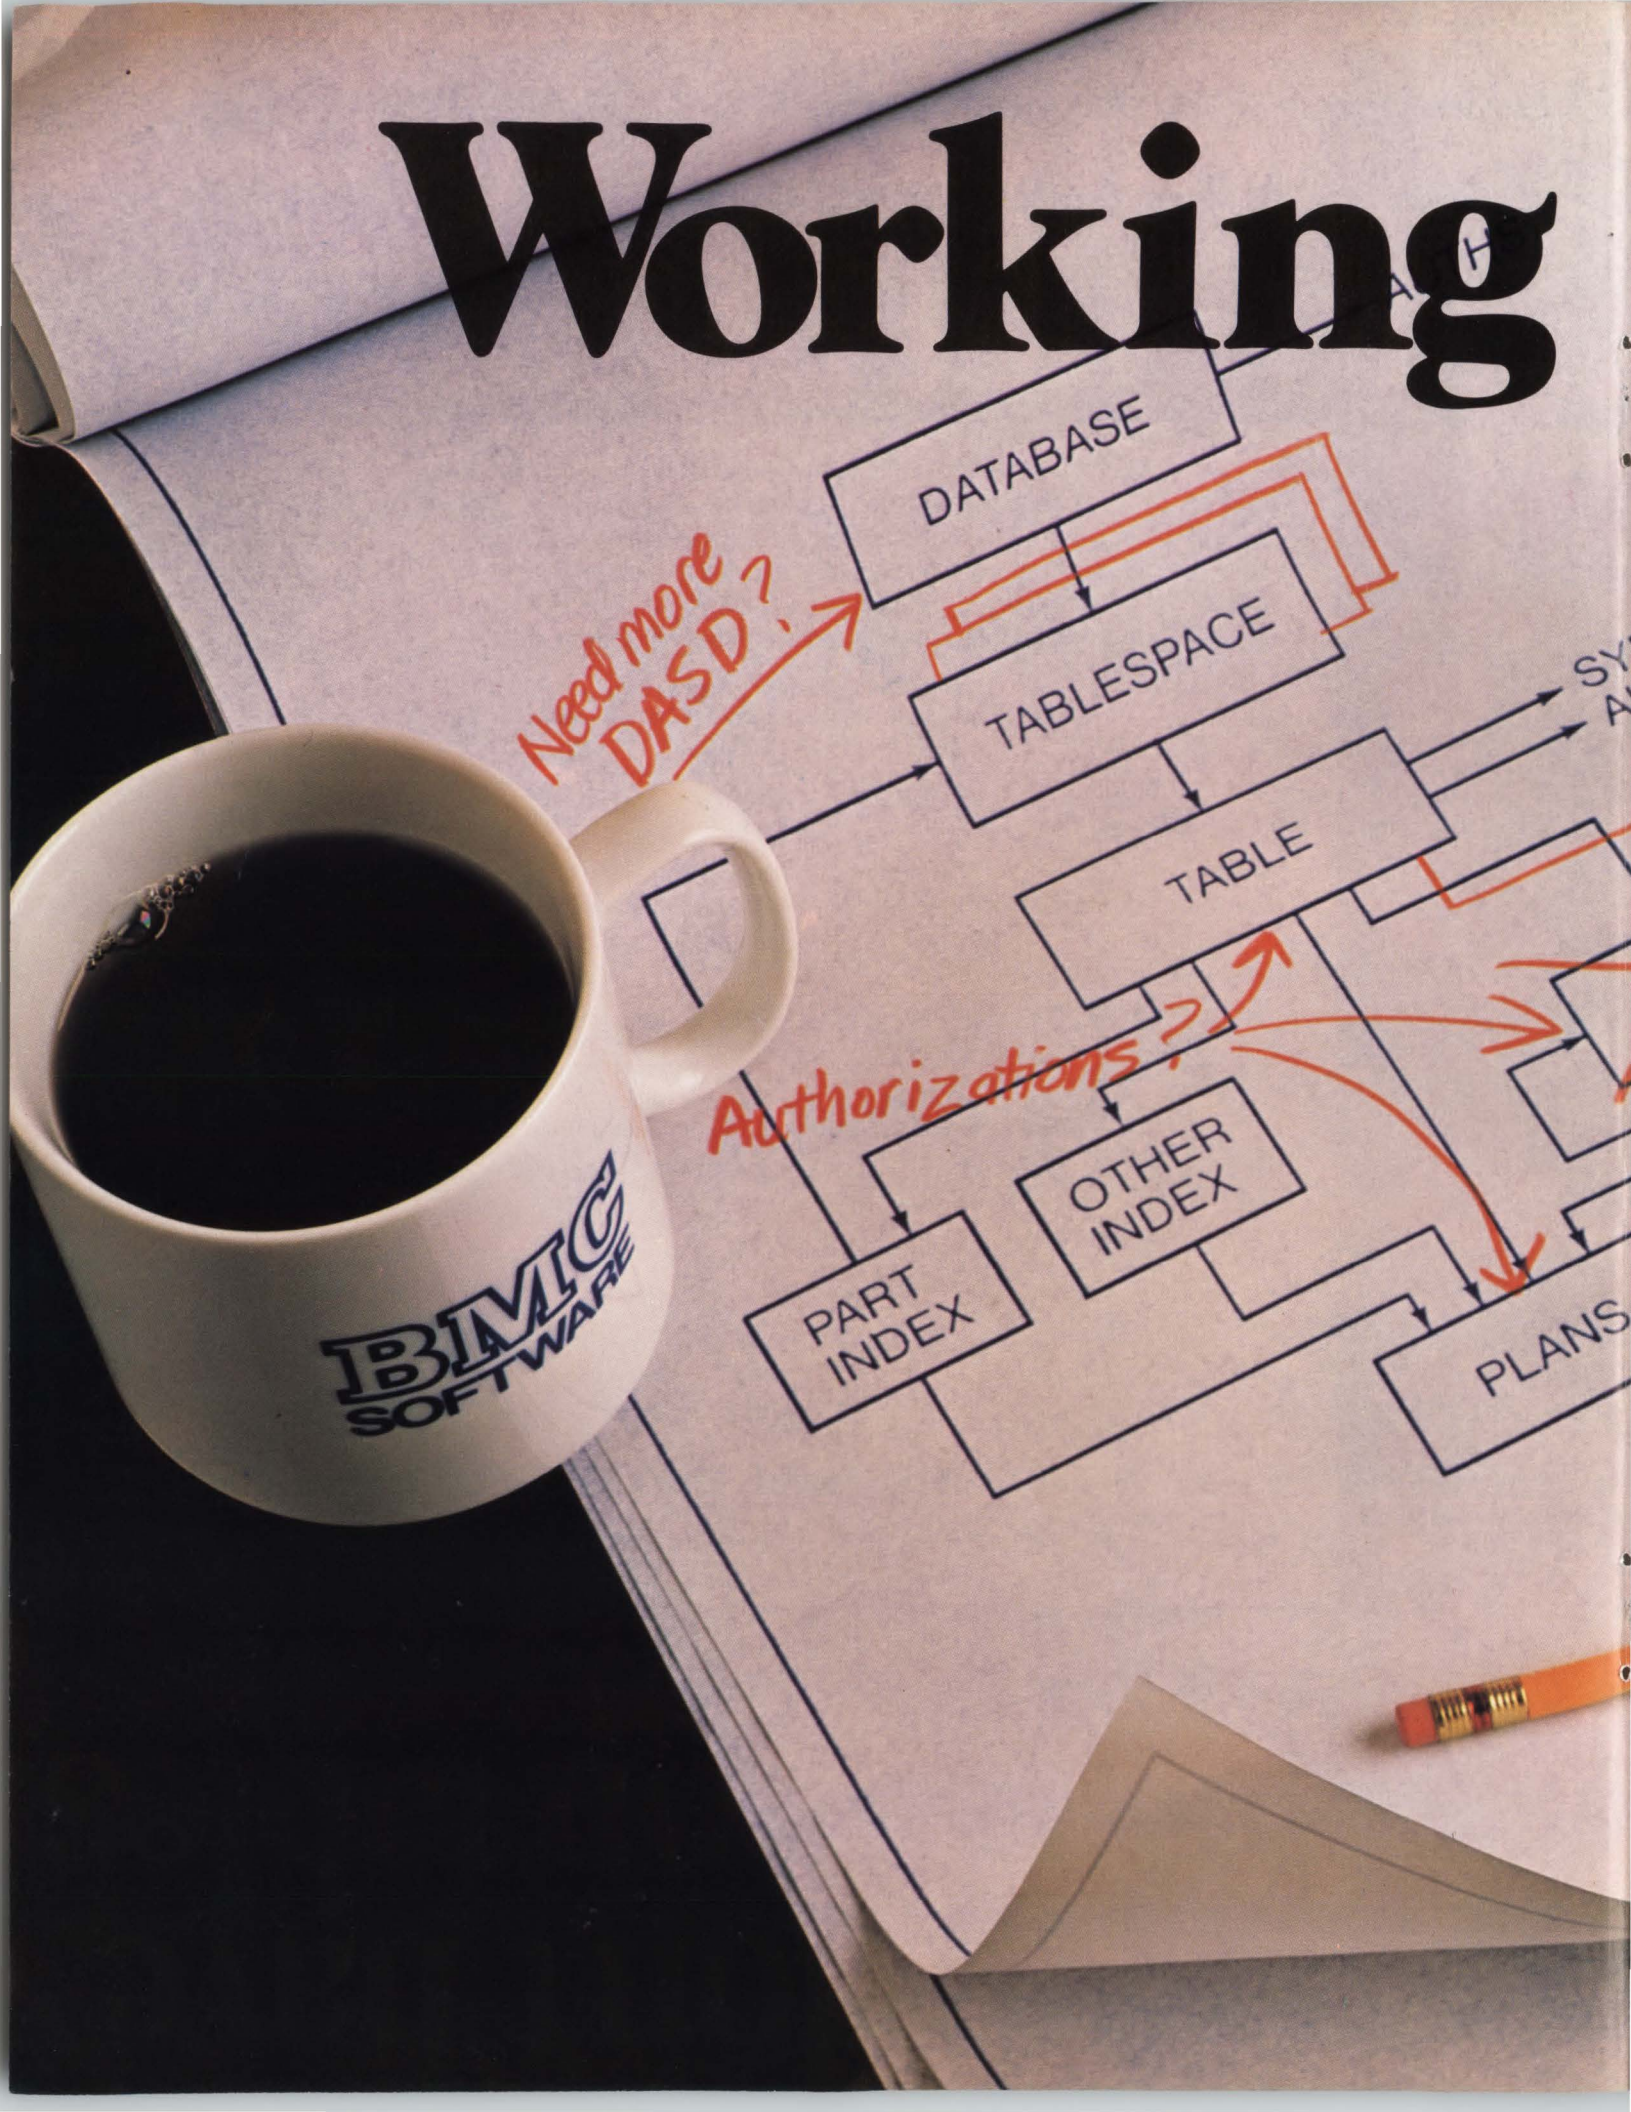
Working
DATABASE
Need more DASD?
TABLESPACE
TABLE
Authorizations?
OTHER INDEX
PART INDEX
PLANS
BMC SOFTWARE

with DB2?

Work with BMC

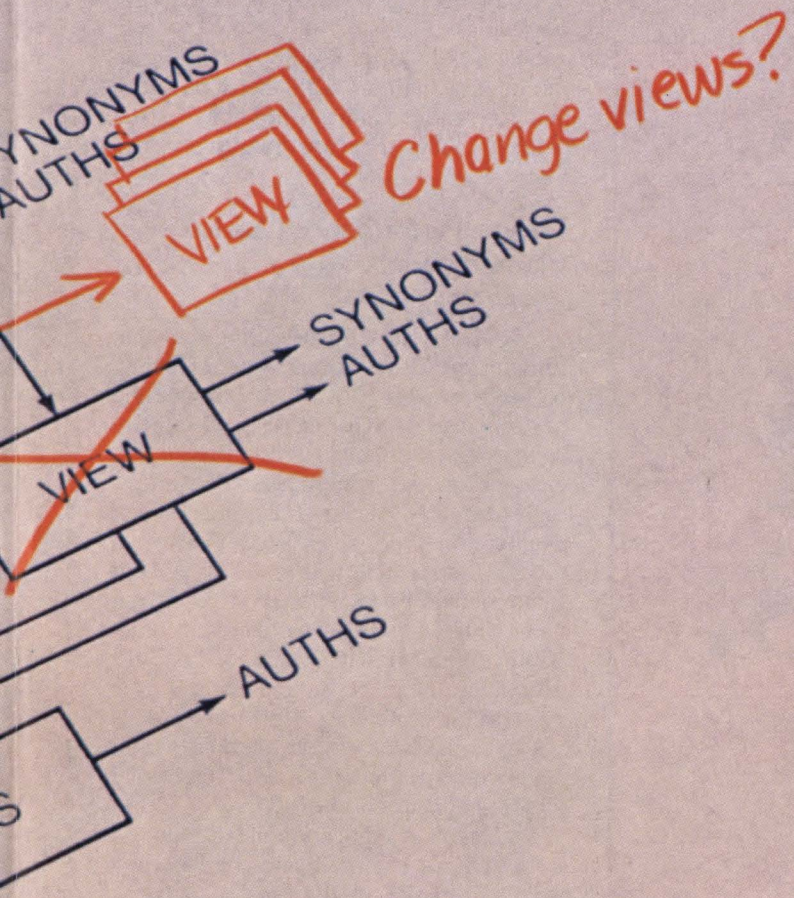

DB2 can be a lot more work than you expected with quite a bit less help than you need. But when you've got BMC Software's comprehensive set of data base administration tools—which include standard interfaces and integrated function—you can reduce your costs and make your work fast, easy and error-free.

DB2 ALTER—provides complete support for changing, copying and migrating DB2 data structures; includes data conversions, authorization-id switching and restart capabilities.

DB2 CATALOG MANAGER—gives quick and easy catalog information, execution of SQL DDL and DB2 utilities, audit logs and extended SQL function.

DB2 DASD MANAGER—controls the life cycle of physical objects with comprehensive space analysis statistics; also includes space estimation, AMS command and utility jobstream generation and action triggers.

DATA PACKER™/DB2—reduces DASD requirements for DB2 tables an average of 50% to 70%; reduces EXCPs.

DB2 REORG PLUS—reorganizes DB2 tables 4-10 times faster than the supplied DB2 utility; provides dual image copy and statistical history.

For more information or to begin a 30-Day-Plus Free Trial of any or all of these products, complete and return this coupon or call BMC Software, Inc., **The Complete DB2 Company.™**

BMC Software, Inc.
P.O. Box 2002
Sugar Land, TX 77487-2002

1-800-841-2031
or call collect: 713-240-8800

BMC SOFTWARE

Contact me about:

Free Trial	More Info	
☐	☐	DB2 ALTER
☐	☐	DB2 CATALOG MANAGER
☐	☐	DB2 DASD MANAGER
☐	☐	DATA PACKER™/DB2
☐	☐	DB2 REORG PLUS
☐	☐	All BMC DB2 products

Name _____

Title _____

Company _____

Address _____

City _____ State/Prov. _____ Zip/P.C. _____

Phone _____

MJ

CIRCLE #103 on Reader Service Card ▲

an overriding decision upon the theoretical philosophies of relational systems design and much of the data will have to be de-normalized back to fourth-or-third-normal form as the tables are designed. The reason for this anomaly is quite simple: performance, system resource consumption and end-user terminal response times *must* be a major consideration throughout the entire design and implementation process. The goal is to reduce the number of physical I/Os necessary to satisfy the various functional processing requirements.

At this point data will be logically and physically re-combined and data redundancies created to reduce the size (number of rows) of the tables, the number of tables and to reduce the number of I/Os necessary to process the SQL requests. Resource consumption and response times will be greatly reduced when a query can be satisfied from one table instead of joining two or more tables.

There are two additional realities of application systems design that will affect the data normalization process. First, application development project cycles and deadlines rarely allow sufficient time for complete and thorough systems and application specifications, let alone extensive time for data normalization. The time necessary for the complete normalization of data for a large application could easily be several months.

Second, if the application designers are thoroughly familiar with the application and its data, significant amounts of the data normalization are intuitively obvious.

Design/Performance/Capacity Planning Tools

ANDB2

ANDB2 is an IBM modeling tool that runs on the IBM HONE system. It may be used by your IBM SE to estimate the performance of SQL queries and utilities against specific tables. While it is not an easy tool to use, comparisons of the timings from ANDB2 to actual queries run against prototype tables have shown that ANDB2 is accurate within the 10 percent figure quoted by IBM. In most cases, the CPU and elapsed times were accurate within 5 percent. The product is good for estimating the response and elapsed time for specific functions. However it does not have the capability to predict overall resource consumption for two, ten or one hundred concurrent users.

The level of application detail required by ANDB2 is extensive and at a very low level. Examples of the type of data required are:

- Number of rows per table, length of rows and number of columns
- Number of indices
- Whether primary index is clustered
- Is access through primary index
- Number of Leaf Pages, number of NPAGES
- First key cardinality, full key cardinality
- Number of rows scanned, number of rows selected
- Number of rows qualified by the RDS and Data Manager
- Is sorting performed, if so, how many rows.

As you can see from this partial list, preparing for the use of ANDB2 can be quite time consuming. It also requires a

YOU JUST LOST ALL YOUR ACCOUNTS.

Your accounting data files could be destroyed at any moment. But there's a quick way to recover: the Journal Processing Environment (JPU/E) from Softsystems. JPU/E is the most powerful and flexible data recovery software on the market– and our journal processing and recovery environment makes it even better. With JPU/E, you can automatically reconstruct VSAM files and DL/I data bases. Automatically archive CICS disk journals. Manage your journal and recovery environment. Create batch journals. And much more. So call Softsystems today. Before you lose all of your accounts.

JPU/E Data Recovery Environment
Because Accidents Will Happen.

SOFTSYSTEMS, INC.

311 Mallick Tower · One Summit Avenue · Fort Worth, Texas 76102
800-331-7846 · 817-877-5070
In Canada: 800-367-8673

Available for VSE, VSE/SP, VS1, MVS, and MVS/XA installations.

CIRCLE #10 on Reader Service Card ▲

IBM mainframe application software has a new DB2 champion

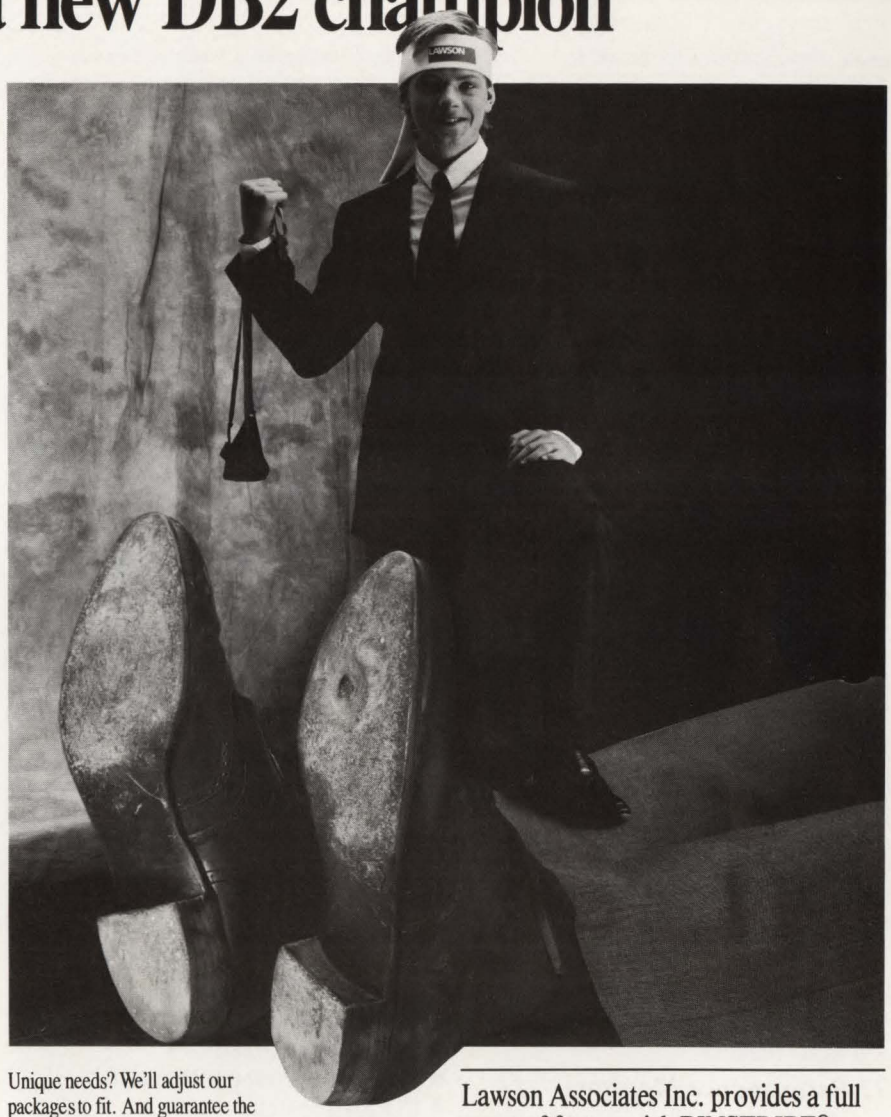

We're Lawson, the IBM mainframe business application software company with the confidence to challenge the industry's Goliaths. And hit them with new ideas. Like low cost, low maintenance—yet full-featured—business application software.

Lawson: Leader in software technology.
Our new technology is rapidly showing the "big guys" a few things about PINSTRIPE® low cost, low maintenance software. Forget "other" IBM mainframe software. With its high prices, old code, time-consuming installation and expensive upkeep.

New technology made our lower costs possible. It also made us first with application software for the DB2 relational DBMS. Now you get a choice: DB2 or VSAM. With newly developed code that's easy to install. Simple and inexpensive to maintain.

Solid application technology.
Comprehensive, well-thought-out application features reflect the experience of our package designers. Seasoned accountants, CPAs and human resource professionals who know what features and functions your users need.

And features they request. In the last four years, 84 percent of user-group suggested enhancements have been added to Lawson PINSTRIPE® packages.

Service and support — cutting the "big guys" off at the knees.
We're big enough to get the job done—but small enough to be sensitive to our clients' needs. With the best service and support available on the IBM mainframe market today. Including your personal software consultant, an expert who knows your needs *and* your application.

A full range of optional Support-Plus services is available, too. Like implementation consulting, data conversion, technical and application consulting. And a post-sale audit to make sure you're utilizing every feature we've built into PINSTRIPE® packages.

Unique needs? We'll adjust our packages to fit. And guarantee the work. If you'd rather do the job yourself, you'll find Lawson software easy to modify. And we'll still support the rest of the package. Other application software companies won't.

Lawson, IBM mainframe software's new champion.
We've put it all together. New technology with fresh code, including both VSAM and DB2 expertise. Full-featured application technology. Backed up by support that sets a new industry standard.

All in low cost, low maintenance application software from Lawson. No wonder we can take on the "big guys." And win.

Lawson Associates Inc. provides a full range of feature-rich PINSTRIPE® business application software packages for the IBM mainframe hardware environment. Available in the U.S., Canada and overseas.

Accounting
• General Ledger • Accounts Payable
• Accounts Receivable • Fixed Assets
• Project Accounting

Human Resources
• Payroll • Personnel • Benefits

Distribution/Procurement Management
• Order Entry • Inventory Control
• Purchase Order

LAWSON

...WHATEVER IT TAKES!

For more information, call
Sue Weinacht
(612) 379-0258
LAWSON ASSOCIATES INC.
1300 Godward Street, Minneapolis, MN 55413

CIRCLE #32 on Reader Service Card ▲

FIGURE 1

```
ACTUAL FROM 10/13/88 07:00:46.53          DB2 PERFORMANCE MONITOR (V1 R0)          DB2 ID:   DB2B          PAGE      1

                                                ACCOUNTING TRACE                    REQUESTED FROM 10/13/88 06:50:00.00
                                                                                              TO   10/13/88 23:30:00.00
                                              BY AUTHID  /PLANNAME

ACCOUNTING TIME      09:34:09.08                      APPLICATION
                                                         TIME          DB2 TIMES
                                                       (CLASS 1)       (CLASS 2)          TERMINATION CONDITION
                                                      -----------     -----------         --------------------------------
AUTHID    -                          ELAPSED TIME / INVOC    1:45.76         1:33.95        NORMAL      DEALLOCATION
CORRNAME  -                          CPU TIME (SSSSSS.TH)      28.86           24.85
CORRNMBR  -  'BLANK'                     TCB TIME              28.75           24.75
CONNECT   -  BATCH                       SRB TIME               0.11            0.10
PLANNAME  -

MANIPULATIVE SQL ACTIVITY   TOTAL                BUFFER MANAGER TOTALS      POOL 0    POOL 1    POOL 2    POOL 32   TOTAL
------------------------- --------              ------------------------- --------  --------  --------  --------  --------
QUERY/UPDATE ACTIVITY                            GET PAGE REQUEST              4915         0         0         0      4915
     SELECT                    27                BUFFERPOOL EXPANSIONS            0         0         0         0         0
     INSERT                    16                SYSTEM PAGE UPDATES             48         0         0         0        48
     UPDATE                     0                SYNCHRONOUS READ I/O          1028         0         0         0      1028
     DELETE                     0                SEQ. PREFETCH REQUEST          240         0         0         0       240
     PREPARE                    0
                                                 LOCKING SUMMARY            TOTALS    SUSPENSIONS (CLASS 3)        ELAPSED TIME
OTHER ACTIVITY                                   ------------------------- --------   ----------------------       ------------
     DESCRIBE                   0                SUSPENSIONS                     0   I/O SUSPENSIONS                   18.16
     OPEN                     106                DEADLOCKS                       0   LOCK/LATCH SUSPENSIONS             0.00
     CLOSE                    106                TIMEOUTS                        0
     FETCH                  10696                LOCK ESCALATION (SHARED)        0
     COMMIT                     3                LOCK ESCALATION (EXCLUS)        0
     ABORT                      0                MAXIMUM PAGE LOCKS HELD         5

CONTROL SQL ACTIVITY        TOTAL
------------------------    ------
LOCK:TABLE                     0
GRANT                          0
REVOKE                         0
INCR. BINDS                    0
COMMENT                        0
LABEL                          0

DEFINITIONAL SQL ACTIVITY   CREATE  DROP  ALTER
------------------------    ------ ------ ------
TABLE                          0      0      0
INDEX                          0      0      0
TABLESPACE                     0      0      0
STORAGE GROUP                  0      0      0
DATABASE                       0      0
SYNONYM                        0      0
VIEW                           0      0
```

detailed level of knowledge both about DB2 and how DB2 will access the tables for each SQL request.

DB2 Performance Monitor

The DB2 Performance Monitor (DB2PM) from IBM is an effective way to analyze the overall performance of DB2 and of individual queries. DB2 provides System Statistics, Application Accounting and several types of performance traces that have the output directed either to SMF (default) or GTF. DB2PM can produce several types of reports at summary, detail and trace levels.

The report type of most benefit to an application in the development process is the Accounting Trace Report. The report itself may be selected by time span and may selectively produce reports for one or more AUTHIDS, PLANS, CORRNAME, CONNECT and so on. Figure 1 illustrates an Accounting Trace Report for one execution of a PLAN by a specific user (AUTHID). The specific items of interest on this report are:

- Number of Selects and Inserts
- Number of Fetches
- Get Page Requests
- Synchronous Read I/O
- Sequential Prefetch Requests
- I/O Suspensions
- Maximum Page Locks Held (this means maximum concurrent, not a count of the number of locks taken while the table was scanned).

The ratio of (Get Page Requests/Read I/O) is approximately 4.8 to 1 indicating that DB2 made little use of its Pre-Fetch capabilities and that there was a low ratio of buffer hits for requested data. (The following average times for a given function are particular to this program/plan and could vary significantly from the averages for a different program/plan.) The Fetch count indicates that more than 10,000 rows of data were returned to the application program. Each Fetch itself required between .00061 & .00074 seconds (and in this program very few Fetches caused I/O). Each overall interaction between the application program and DB2 to Fetch the qualified data into the program averages .01 seconds times 10696 Fetches which is equal to 106.96 seconds of the total elapsed time. The tables were not "opened" prior to the start of the run and this accounted for an additional 11.36 seconds. The 16 Inserts averaged .47 seconds each or 7.52 seconds. The 27 Selects averaged .26 seconds each or 7.02 seconds. To summarize this data:

Fetch Time	106.96
Opening Tables	11.36
Inserts	7.52
Selects	7.02
Total	132.85 Seconds

The balance of the time would be attributable to opening and closing Cursors, Commits, loading the SKCTs and DBDs into the EDM Pool and the Create/Terminate Thread functions.

If it appears that a specific user function or PLAN has a performance problem such as excessive numbers of page locks, lock suspensions or escalations, deadlocks or unexplainable elapsed times, various types of special performance traces may be used to isolate the causes of the

COMPUWARE CICS dBUG-AID

Find Program Errorrs Fast

If you didn't spot our errorr Compuware CICS dBUG-AID would. This versatile debugging and testing tool helps detect and isolate multiple program errors, logic errors, CICS transaction abends, even CICS storage violations. Fast. Because with CICS dBUG-AID, programmers can interactively view text one line of code at a time. And, its unique three-tier design enables the new user, as well as the more experienced programmer and systems programmer, to put CICS dBUG-AID to work immediately. With CICS dBUG-AID you can't go wrong. To find out more about CICS dBUG-AID, CICS Abend-AID, and CICS PLAYBACK, call or write Compuware at: 31440 Northwestern Hwy., Farmington Hills, MI 48018-5550, 1-800-521-9353 or (313) 737-7300.

COMPUWARE

Because experience is everything.

Compuware CICS dBUG-AID and CICS Abend-AID are registered trademarks of Compuware Corp. Compuware CICS PLAYBACK is a trademark of Compuware Corp. © 1988.

CIRCLE #67 on Reader Service Card ▲

FIGURE 2

DB2 - Table Retrievals via Clustered Index

Job Name	Elapsed Secs.	CPU Secs.	SRB Secs.	# Rows Rtrnd.	# Rows Sel.	# Fetches	# Get Pages
1. Test1	117	90.53	.20	58820	70	58890	4587
2. Test2	168	101.55	.78	79220	350	79570	4589
3. Test3	478	283.89	2.32	250152	1155	251307	7235
4. Test4	75	37.52	.29	22680	88	22768	7849
5. Test5	228	116.51	.88	79220	70	79290	18134
6. Test6	511	290.80	2.29	250152	165	250317	14656

Notes:

- **The table contained 1,244,878 rows.**
- **Access through a Clustered Index.**
- **Non-Sargable predicate which forced final selection criteria in the RDS.**
- **Note that the elapsed times and resource consumption comparing Jobs 1-3 and 4-6 are somewhat less than linear.**

HANDS-ON DB2

You've decided to go with DB2 – and now the deadlines are looming! Your staff may be new to DB2, but they have a wealth of experience and talent to build on. What they're saying is "Give me the background, show me how it works and how it fits in our environment – that's what I need". And that's what we deliver – effective, thoroughly tested, hands-on training.

DB2-VERSION 2 IN ONE WEEK Reserve your class now.

Data Kinetics
Education Division
Tel:(613)238-6709
Fax:(613)238-2852

HANDS-ON WORKSHOPS FOR:
DB2 ■ IMS ■ FOCUS ■ VM/CMS ■ MARK IV
tableBASE ■ QNX ■ ANSWER/DB

CIRCLE #175 on Reader Service Card ▲

problem. These traces must be used judiciously, as the improper use can add 100 percent of CPU overhead to the DB2 system, cause many thousands of records to be written to SMF and even cause SMF to fill its datasets and abend. Aside from the catastrophic effects this can have for the entire system, the volume of records produced can make the analysis range from difficult to impossible. Looking at a few hundred pages of trace information is difficult; looking at many thousands of pages is just not possible.

CRYSTAL/DB2

CRYSTAL is a product from BGS Systems, Inc. (Waltham, MA) and provides a facility for modeling application systems. To the best of my knowledge, the CRYSTAL software is the only facility available that allows the creation of detailed application models for IMS, CICS and DB2.

I have used the CRYSTAL/IMS software for modeling large IMS/VS-based applications for several years. The software provides data about individual transaction/workload resource consumption and response time, overall CPU consumption and DASD utilization. It is quite useful in selecting/sizing a processor and planning for DASD by highlighting any potential bottlenecks due to having several high activity databases on the same physical volume. It may also indicate the need to split (partition) databases and distribute the partitions across multiple physical volumes to reduce device contention.

Again, a great deal of detailed application knowledge is necessary to build an accurate model. The accuracy of the CRYSTAL model will be directly proportional to how accurately the modeled application represents the real application. In most cases, the modeling should be an iterative process by building a base model and continually modifying it as the application code is developed. CRYSTAL provides both on-line and batch modeling facilities. Both of these facilities may be used to project various "what if" scenarios during one iterative execution of the model. A common example might be to project CPU utilization and transaction response times at transaction rates of 5,000, 7,500 and 10,000 per hour.

IMS DC MONITOR (if IMS Attach)

The IMS/VS DC MONITOR is useful when IMS Attach is being used to provide the on-line transaction management for

BIG BANG

If that's what you've heard about mass conversion, forget it. That's just our competition making noise because CTG/CORTEX is light-years ahead of every other conversion technology. A 100% success rate proves that mass conversion is the safest approach. For DOS to MVS on a fixed schedule, for a fixed price, call 1-800-DOS 2 MVS.

CTG/CORTEX

COMPUTER TASK GROUP, INC.
3095 Union Road • Orchard Park, New York 14127
(716) 674-9310 • 1-800-DOS 2 MVS

CIRCLE #26 on Reader Service Card ▲

FIGURE 3

DB2 - Sort Time Comparisons

Select - Tablespace Scan - Batch Runs With/Without Sort (order by)

Test Example	Rows Returned	Columns Returned	Elapsed Seconds	CPU Seconds
Run 1 w/o SORT	225568	4	784.0	349.30
Run 2 w/o SORT	439782	4	878.0	526.47
Run 3 w/o SORT	655424	4	1255.0	906.60
Run 1 with SORT	225568	4	1204.0	411.04
Run 2 with SORT	439782	4	1484.0	704.15
Run 3 with SORT	655424	4	1821.0	1123.06

Notes:

- **The test jobs were run on a 3081KX, lightly loaded with batch work after 7 p.m.**
- **The table has 1,244,878 rows.**
- **The table has one index, first key cardinality is 1, and full key cardinality is 995.**

FIGURE 4

DB2 - Column Retrieval Time Comparisons

Select - Tablespace Scan - Batch Runs, No Sort

Test Example	Rows Returned	Columns Returned	Elapsed Seconds	CPU Seconds
Job 1	134474	2	230.0	179.73
Job 2	134474	4	255.0	194.27
Job 3	134474	6	271.0	209.38
Job 4	204476	0	145.0	43.46
Job 5	204476	2	307.0	262.53
Job 6	204476	4	327.0	277.38
Job 7	204476	6	366.0	296.74

Notes:

- **The test jobs were run on a 3081KX, each job was executed four times and the best (lowest) times are illustrated. The wall elapsed time variations are not of any real concern. The CPU elapsed times within each set of jobs was within 1.5 seconds.**
- **The table has 1,244,878 rows.**
- **The table has one index, first key cardinality is 1 and full key cardinality is 995. This was not used for the tablespace scan.**

DB2/SQL requests. The DC MONITOR does not provide any actual data about what has occurred in DB2. It does provide the elapsed time for the request to be passed to DB2 and processed. This, in conjunction with DB2PM reports, would highlight any performance problems in the attach mechanism and two-phased commit process.

A Design Experience . . .

One particular design and prototyping exercise is especially valuable. The application deals with large volumes of data that have high insert, delete and update requirements.

Based upon the data normalization processes, the primary application table was originally projected to contain 150 million unique rows. One-third of this volume would be both added and deleted on a weekly basis (delete 50 million rows of old data and add 50 million of new data).

It rapidly became apparent that a table of this size was not a viable approach due to the magnitude of the update/delete activity, even though most of this activity would not be on-line. This volume of activity, in addition to lengthy batch update run times, would create problems due to archiving of the log datasets and requiring frequent reorganizations. After the large update and reorganization runs, it would be necessary to execute 'RUNSTATS' to update the system statistics for the table and to re-bind all the plans to ensure efficient table accesses.

Subsequent de-normalization of data resulted in seven application-partitioned tables, each containing approximately two million rows of data. This provided greatly reduced run times for all update, query and utility requirements.

The extensive testing and monitoring of the application during the development process provided much data that will be useful during the development of other applications. Some examples are as follows.

- The elapsed and CPU times to retrieve data through a clustered index, as percentages or multiples, is slightly less than linear. For example: using the same table that contains 1.2 million rows query 'A' returns 4,000 rows and query 'B' returns 40,000 rows, query 'B' may take approximately 7-8 times as long. This is illustrated in Figure 2. The numbers are slightly different; however, the ratios and effects upon elapsed and CPU times are representative.
- When DB2 has to perform sorting, the elapsed time for the retrieval will increase significantly as illustrated in Figure 3. Sorting will be caused by joins of two or more tables, ordering the results of a query and so on. If large numbers of rows are being retrieved from large tables, an application should consider the usage of an external sort.
- Joining of tables will cause the creation of temporary work tables, generate "system page updates" and will drive

See DB2 *page 103*

EXCLUSIVE FROM MILTOPE!

Continuous Fan-Fold Printing at 37/90 Pages per Minute

Cut sheet printers available at 30/60/75 ppm

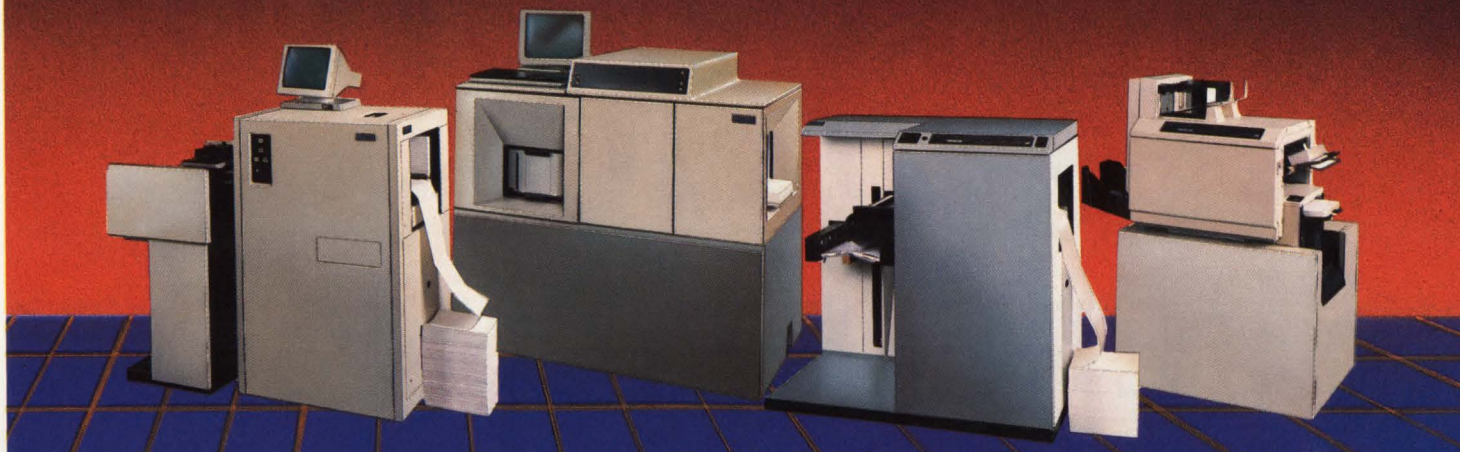

Miltope's FAMILY of Non-Impact 'ion' deposition page printers affords letter quality printing at 'less than a penny' per page.

Miltope provides multifont alpha/numeric printout, ''true'' electronic forms overlay plus sophisticated graphics for generation of signatures, logos, bar codes, OCR and special characters.

Miltope's page printers eliminate time-consuming bottlenecks by combining speed, high-duty cycle, reliability and versatility in electronic printing systems that are plug-to-plug compatible and fully emulate:

- Xerox 3700/2700
- HP LaserJet Plus
- HPGL
- HASP
- QMS® MAGNUM® (Printronix Graphics)
- IBM Interfaces: 3211 — 3287 — System 3X — AS/400 2780/3780 — 3270 — 3770 etc.
- Dataproducts or Centronics Interface
- Mini/Super-Mini Interface (i.e., DEC, DG, Wang, Unisys, etc.)

The choice is YOURS!

MODEL	SPEED	FORMAT	RESOLUTION	MONTHLY VOLUME
3801	90 PPM	Continuous Form	240 × 240 DPI	Over 1 Million Pages
SERIES 75	75 PPM	Cut Sheet	240 × 240 DPI	Up to 1 Million Pages
SERIES 60	60 PPM	Cut Sheet	240 × 240 DPI	Up to 1 Million Pages
SERIES 37	37 PPM	Continuous Form	300 × 300 DPI	Over 250,000 Pages
SERIES 30	30 PPM	Cut Sheet	300 × 300 DPI	Up to 250,000 Pages

Cost Effective Printing and Dependable Nationwide Field Service have made Miltope the ''Source'' for ALL ion deposition printing systems.

''IMAGINE YOUR IMAGE''

MILTOPE **Business Products, Inc.**

1770 WALT WHITMAN ROAD • MELVILLE, NY 11747

TEL: 516-756-7650 • TWX: 510-221-1803 • FAX 516-756-7606

Miltope Printers use Delphax ion deposition engines.

QMS and MAGNUM are registered trademarks of QMS, Inc.

CIRCLE #27 on Reader Service Card ▲

STRATEGIES FOR IMPROVED SQL/DS PERFORMANCE

The flexibility of SQL/DS gives unaware users just enough rope to hang themselves.

By Robert D. Smith and Lani Spund, PhD

Interest in SQL/DS is on the upswing again. There are a number of reasons for this revival but chief among them is the dawning awareness that SQL/DS, introduced by IBM in 1983, is an incomplete tool for realizing the main benefits of relational technology. To gain all the benefits, the DP shops encountering SQL/DS for the first time will have to change their mindsets in terms of how they develop systems. The cultural changes required by SQL/DS and the relational model is the subject of this article.

If you are the MIS manager in the medium-sized DOS/VSE shop who may be the first to look at SQL/DS, you will have to fundamentally modify the traditional approach to systems development or risk missing the benefits of the Relational Database Management System (RDBMS). If data modeling is something you have not paid much attention to, you will. You will also learn not to overdo it. A little bit of data modeling goes a long way in improving design. Moreover, you will have to pay attention to a number of issues that were, at best, of peripheral interest before.

One issue is that SQL/DS is a retrieval-based data manipulation sub-language for users in the IBM DOS/VSE and VM environments. Its big brother product, DB2, is designed to run under MVS. Generally used by relatively small DP shops, SQL/DS gives smaller shops access to the advantages of RDBMS applications.

Another issue is that SQL/DS is an important enabling technology to implement relational applications. It is lacking three components, however, that are essential for modern organizations wishing to develop and manage mission-critical applications. The three components are: an active data dictionary, integrated Fourth-Generation application development Language (4GL) and data security. When these components are all integrated with an enabling technology like SQL/DS and a set of sophisticated management processes, a user has an extremely powerful platform for application development and management.

Benefits of SQL/DS

Although SQL/DS has its limitations, it is a well-accepted tool with six benefits: relational applications, data integrity, integrated data dictionary, isolation of applications from Database Management System (DBMS) structure, programmer productivity and security.

Relational Applications

RDBMSs have emerged from the antiseptic laboratory of the academic world and have joined their traditional counterparts in the down and dirty world of production processing. Moreover, an increasing number of companies have exploited the technology in support of mission-critical, front-office applications.

The chief benefit of the relational model is not in its physical implementation but in the way it causes data to be viewed, accessed and presented. The relational model simplifies the view of data and clarifies the logical associations in an understandable and common-sense manner. This simplification allows data to be accessed and utilized by a far broader class

of users than the traditional structures that are noted for their rigidity.

The resulting flexibility is now allowing companies to do more with their strategic information assets than merely running applications like inventory or producing standard reports. With much more dynamic access to data, companies can now gain a much better understanding of their day-to-day operations and also see both short-term tactical and long-term strategic advantages over their competitors.

Data Integrity

Mission-critical applications place stringent requirements on the underlying software. For such applications, companies are insisting on data integrity in their DBMSs. In a distributed database environment, records must be updated in a consistent, thorough manner in order to avoid data corruption. A corporation's data assets must be synchronized to ensure that all employees are using the same values.

Another form of data integrity is referential integrity. Referential integrity is an ugly name for business rule validation. A database has referential integrity if it "hangs together." Every personnel record, for example, refers to a real employee in a real department. Every part number refers to a real part record. Every work order refers to a real customer. It is easy to enforce simple business rules (that is, a decimal number field must not have alphabetic characters). But integrity within and across applications is trickier. For example, how do you ensure that a customer's record is not deleted if there are still outstanding invoices for that customer? Database vendors are beginning to help but most analysts agree that the best way to solve this problem is to build referential integrity right into the database engine instead of into individual applications.

Integrated Data Dictionary

The trend toward installation of non-central DP sites has made the data dictionary more important than ever. Unfortunately, the absence of an integrated, active data dictionary in SQL/DS has seriously hampered product development and acceptance. A good data dictionary offers a powerful and flexible cross-referencing facility that documents the data sources independently of hardware, software or organizational structure. The totality of interactions between systems and databases constitutes the information environment of the enterprise. The data dictionary acts as an automated tool for standardizing and controlling the definitions of and relationships between the processes (systems, programs, modules) and information (databases, files, individual data items).

One of the primary benefits of using SQL/DS integrated with an active data dictionary is optimization of data sharing through minimization of unplanned redundancies.

Since it is so easy to duplicate SQL/DS systems, users may be tempted to needlessly proliferate applications and create unnecessary maintenance headaches. Ironically, this is not a problem under traditional development methods because of their lack of integration that makes it unlikely that users will reinvent something that already exists.

Isolation of Applications from DBMS Structure

Applications built with SQL/DS offer a much more complete separation of ap-

Immediate Corrective Action

Detect and correct mismatches online as they occur. Bypass, resynchronise or correct in just one session. Eliminate multiple batches and runs. No new scripts. No sequencing errors. *Immediate corrective action.* ▶ CICS/REPLAY lets you review, perfect and execute CICS applications and releases online automatically in just one session. You get a clean cut over into production—and a single person can do it all. ▶ Find out more about how CICS/REPLAY keeps costs down and systems up. Call: Interactive Solutions, Inc. 201-488-3708/212-475-2400

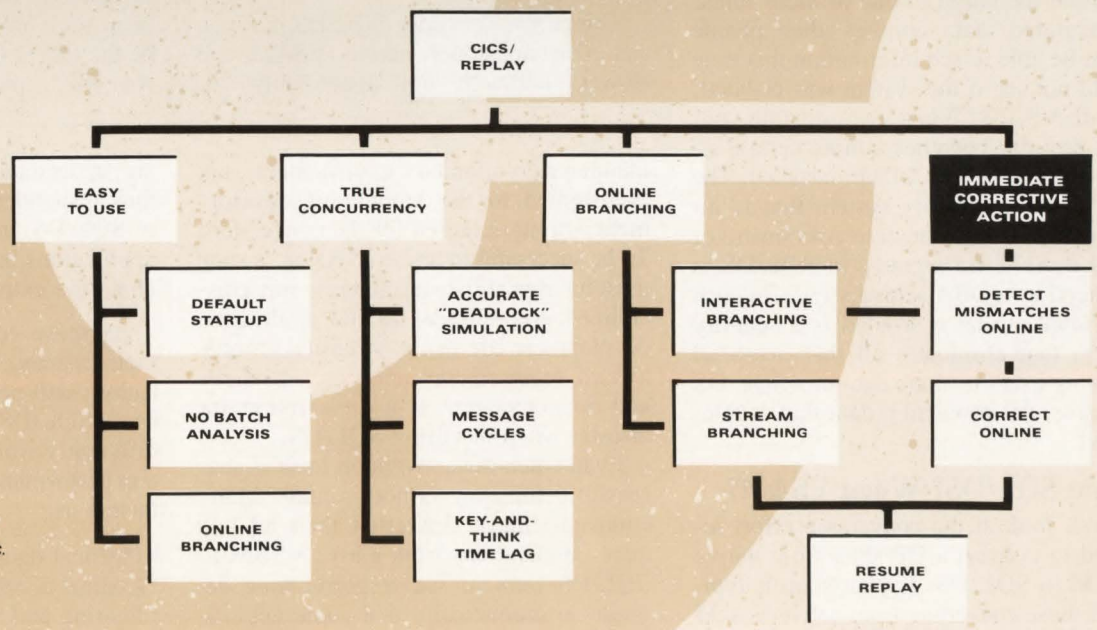

INTERACTIVE SOLUTIONS, INC.
A TBS Company
53 W. Fort Lee Road
Bogota, NJ 07603-1202

plication code from data than was previously possible. The relational architecture allows a company's information assets to be put into a number of physical file structures to suit the performance needs of their various owners. The DBMS engine itself implements the relational model at the conceptual level. This separation of application code from data gives tremendous freedom to developers and users at the all-important logical, user-view level.

Programmer Productivity

Many users find SQL/DS so rigid and restrictive that they clamor for a more complete application development environment. A 4GL application development environment can be a boon to VSE/SP data centers implementing SQL/DS because it reduces the skill sets needed to implement applications. Modest development shops unused to demanding development schedules will find a 4GL represents substantial labor savings when building, maintaining and managing their systems' support portfolios. Examples of 4GL application development environments are CYGNET/SQL (Phoenix Software, Los Angeles, CA), Gener/OL (Pansophic Systems, Oak Brook, IL), NATURAL (Software AG of North America, Reston, VA) and IDEAL (Applied Data Research, Princeton, NJ).

Security

Security becomes extra important with an RDBMS because of the required data redundancy. Every data item that is a Join must be carried redundantly in a Row and Column. Because of that, sensitive data might be the target of one of these Joins.

Integrated data implies that people might be able to see information that they would not see if the system was isolated, as with VSAM. When you integrate your data, security becomes a more central issue regardless of the DBMS installed. The need is for a security system that is independent of the Database Administrator (DBA). As it is, security in SQL/DS is managed by a DBA setting views for each individual. What is needed is a security system that eliminates all that overhead allowing users to share data in a SQL/DS database, yet protecting data that is integrated.

When SQL/DS? When VSAM?

First, look at the conversion effort required to convert a DP shop from native VSAM to SQL/DS. If your shop is typical of those converting from native VSAM to SQL/DS, it is probably in a small maintenance-oriented environment not accustomed to the amount of development activity required for the conversion. To be successful in any conversion, a shop must have established standards and a disciplined commitment to the application development life cycle. A casual attitude toward requirements analysis, standards and documentation is a prescription for disaster when installing SQL/DS.

Even when the conversion itself is successful, efficiency is not assured. Many companies have converted their ad hoc query applications from native VSAM to SQL/DS only to see response time degrade so dramatically that some applications leave skid marks. System managers are understandably confused. After all, they responded to the publicized benefits of SQL/DS and the promise of relational applications but they never thought much about the extra overhead.

As these companies have learned, to their dismay, some categories of applications will perform relatively poorly because SQL/DS implements many features differently from VSAM. In general, system performance is a combination of such factors as:

- Physical data storage techniques
- Logical access mechanisms and indexes
- Sharing and integrity features (locking)
- Features used for recovery (logging)

Five Misconceptions About SQL/DS . . . Or Are They?

1. SQL/DS is DB2 by another name.

The fact is, SQL/DS and DB2 evolved separately and were, in fact, developed in different parts of the world: SQL/DS in Switzerland and DB2 in Santa Teresa, CA. They just happen to have SQL as a common data manipulation language. Given identical statements, what they do under the covers is not necessarily the same. The performance and functions that you encounter in DB2 cannot be assumed in SQL/DS or vice versa.

2. SQL/DS cannot be used to develop full applications.

Of course not. By the same reasoning, COBOL is not a complete development language either. To turn a procedural language into a development language, it is necessary to add a screen generator, a report writer, a debugger or a 4GL. SQL/DS is a retrieval-based data manipulation sub-language that gives users a standardized way to access data stored in RDBMs. SQL/DS benefits greatly from application development tools like IBM's own Cross System Product (CSP) or SAS Institute's (Cary, NC) SAS statistical/graphics package.

3. SQL/DS is inherently slower than VSAM.

This is, in fact, not a misconception. SQL/DS is inherently slower than native VSAM because it represents a higher level of data independence and abstraction. As a consequence, users and programmers cannot take advantage of the physical structure of the database to make their programs run faster which forces slower processing.

Things you do under SQL/DS that you can also do under native VSAM run slower, although not appreciably so. Where SQL/DS tends to slow down is when you do advanced data manipulations that you would not normally do under VSAM because it would be too difficult to do. SQL/DS gives you the ability to do things you could not easily do with VSAM but you pay a performance penalty for having SQL/DS do much of the work for you. In DP shops without an active database, a lot of time is spent sorting and merging.

4. SQL/DS and the relational model is a panacea for my company's data management problems.

Wouldn't it be great! The fact is that SQL/DS is only a tool that allows you to build an information model. It is like a shovel; finding where to dig is the hard part. That requires a management process.

Ever since the DP industry began about 35 years ago, it has been searching for the magic bullet that would solve all problems. For a while it was parallel processing. Later it was Local Area Networks. Now for some people, RDBMs fill that role. Everything gets oversold. SQL/DS is an enabling tool that can be of considerable benefit when supported by a commitment from the enterprise to change the mindset of the system development staff.

5. SQL/DS can never be used directly by end users.

This is true, but not necessarily because the use of SQL/DS requires unacceptable levels of training. SQL/DS is only a way to get at records. It does not display them or manipulate them in any other way. In this way users need a set of development tools in which SQL/DS happens to be the driver that goes and gets the data. It is only a piece of the puzzle.

SQL/DS and DB2: The Differences

IBM's DB2, running under MVS, shares some commonalities with its sister product, SQL/DS, but it is different from SQL/DS in some fundamental ways. First, some similarities.

Like SQL/DS, DB2 arranges its processing components into layers: the Application Program/Interactive User, Relational Data Services, Data Manager and lower level components related to physical I/O. Both the SQL/DS and DB2 systems are buffered, which means that pages for data, indexes, logging and mapping are kept in virtual memory. Additional CPU resources are consumed managing buffer queues that were not used recently and performing buffer paging I/O.

The methods or techniques used by SQL/DS to retrieve data are called paths. For most tables, SQL/DS can follow several paths to find the data requested by a query. The most common path (it exists for every table) is a sequential scan of all active data tables in the DBSPACE. Path selection in DB2 is similar to SQL/DS. DB2 bases its path selection on catalog statistics that are updated only by running the utility RUNSTATS.

But DB2 and SQL/DS handle predicates differently. The path that SQL/DS chooses for each table is determined by the predicates used in the query. Predicates are the conditions specified in the WHERE clause. SQL/DS recognizes three types of predicates: key domain, sargable and residual. A predicate is classified as a key domain predicate if it can be used as a key to an index. "Sargable" is a contraction of the term "search argument." A predicate is sargable if it can be used by the lowest level component of SQL/DS to test rows as they are read from the database. Finally, residual predicates are those that are not classified as key domain or sargable. Unlike SQL/DS, the predicates IN, LIKE and NOT can be sargable under DB2.

Under DB2, unlike SQL/DS, an EXPLAIN can be requested during BIND time allowing a DBA to review all the queries used in an application.

DB2 supports high rates of sequential retrieval and table loading. As a result, the elapsed time for deleting a significant portion of a row from a large table may be reduced by selecting the rows to retain, writing the output to a file and replacing the table from the selected rows. The best information available on DB2 path selection is IBM publication GG24-3004, *IBM Database 2 Application Design and Tuning Guide*.

- Security features
- Memory management operations.

Although VSAM is less "full-featured" than SQL/DS, a VSAM access from a COBOL program requires fewer instructions than an equivalent SQL/DS access. SQL/DS has a longer access path than VSAM because the latter does not automatically implement the sharing, recovery and security functions. Moreover, SQL/DS may use a less efficient physical storage technique and the SQL/DS optimizer may locate and access the data using a less efficient logical technique.

The SQL Standard

IBM's Structured Query Language (SQL) has more than its share of critics. Like the Holy Roman Empire that a wit once noted was neither holy, Roman, nor an empire, SQL is neither structured, query-like, nor much of a language. It is instead a standardized access method to repositories of data. Nevertheless, SQL has emerged as the de facto standard that promises to interface all DBMSs across all vendors.

See SQL/DS *page 98*

Concurrency You Can Count On

Message cycles present messages to CICS at precisely the same instant—or in the original time order. "Key-and-think lags" mimic human operators. Deadlocks identified prior to production. They all add up to *true automatic concurrency testing*. ▶ CICS/REPLAY lets you review, perfect and execute CICS applications and releases online automatically in just one session. You get a clean cut over into production—and a single person can do it all. ▶ Find out more about how CICS/REPLAY keeps costs down and systems up. Call: Interactive Solutions, Inc. 201-488-3708/212-475-2400

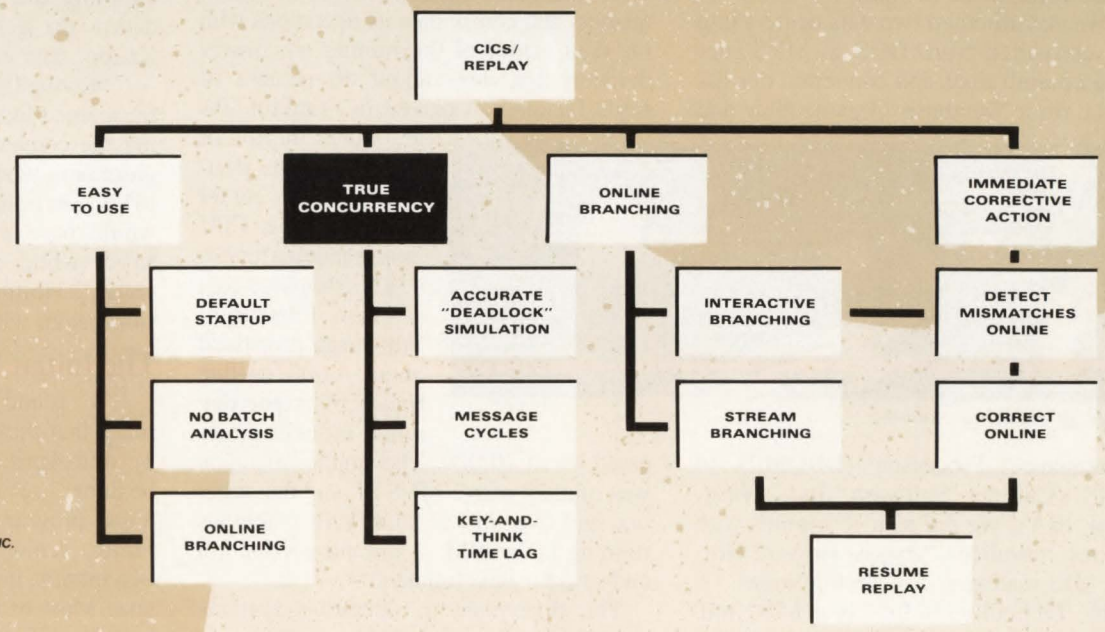

INTERACTIVE SOLUTIONS, INC.
A TBS Company
53 W. Fort Lee Road
Bogota, NJ 07603-1202

Joy Technologies Rides Out Triple Conversion

By Lamont Wood

You can move a data center. Or you can convert to a new operating system. Or you can convert to a new computer. And you might emerge unscathed.

To do all three *at the same time* would seem to tempt fate. Yet that is what they did at Joy Technologies. Not only did the company emerge sane and healthy, but also it saved money and improved response time in the bargain.

"We consolidated two data centers into one, converted from DOS to MVS, upgraded mainframes and converted our database from Database Organization and Maintenance Processor (DBOMP) to IDMS (Cullinet Software, Inc., Westwood, MA). We did it in 15 months with in-house resources," boasts Howard Horton, MIS manager, Joy Technologies.

Longwall shearing machine.

Joy Technologies Inc. is a $500 million, closely-held manufacturer of underground mining equipment headquartered in Pittsburgh, PA. In business since 1924, the 4,500-employee firm decided to divest some peripheral divisions after a merger in 1987. This downsizing led to the 15-month odyssey of the 22-member IS department of Joy's mining machinery division.

That is because, as part of the divestiture, it was decided in January, 1987 to close the data center in corporate headquarters and consolidate its operations with the data center of the mining machinery division that does about 60 percent of Joy's business. Located in Franklin, PA 75 miles north of Pittsburgh, the division's data center boasted a larger CPU than headquarters: a 4381, compared to a 4341 in Pittsburgh. Pittsburgh processed about 4,000 on-line transactions per day while Franklin processed about 70,000. The smaller machine was running native MVS/SP and the larger one had two logical DOS/VSE machines running under VM — one for production and one for development.

The difference in operating systems proved significant and led to a decision to physically relocate the corporate system to Franklin without change until a conversion/consolidation plan could be worked out.

"Here in Franklin, we had been postponing a DOS-to-MVS conversion for some time. Being asked to take on the responsibility for corporate-wide data processing support was a catalyst for our decision to do this conversion before attempting the corporate/division consolidation on a single CPU," recalls Ken Knapp, data center manager.

"Moving the Pittsburgh CPU to Franklin seemed like a convenient way of making the conversion since the system programmers would not have to generate an MVS environment from scratch. They would have a machine already running MVS while converting the DOS machine," Horton explains. So the conversion began with the move of a computer.

The Initial Move

"We found out about the first of February that they wanted to move the 4341 by mid-April. No one believed it could be done," recalls Lou Northrup, lead systems programmer. But he and Larry Smith, senior technical specialist, spent two months in Pittsburgh learning the system. Most of the Pittsburgh staff had already been let go in the downsizing. None

were brought to Franklin and no outside contract programmers were used during the conversion.

"The local control units in Pittsburgh had to be converted to remote units so that the terminals in Pittsburgh could be controlled from Franklin. The remote controllers were actually first installed in Pittsburgh and tested with modem eliminators," Smith recalls. They took photographs of two large data communications cabinets, patch panels and other connections in order to duplicate them exactly at the Franklin site.

Back at Franklin, the biggest problem (other than the fact that no staff additions were made to deal with the new arrival) was putting everything in the 50x80-foot computer room. It was already occupied by the 4381 and a large CAD/CAM system. A power distribution unit was brought from Pittsburgh with the 4341 and all its peripheral equipment. With planning, the space proved adequate, but cramped. Luckily, the air conditioning and Halon fire suppression system at Franklin was already adequate.

"We shut the Pittsburgh machine down on Friday and had it running in Franklin on Monday," Northrup recalls in amazement.

"Corporate users were calling us on Monday to see if we had really moved, because they could not tell any difference. The move seems so simple after the fact, but at the time we were filled with concerns," Knapp comments.

Software Conversion

The switch to MVS then got underway in earnest. Technical staff members were sent to MVS training for the next couple of months. Converting to MVS meant abandoning the familiar utilities and applications of DOS. "In most cases they were very different and generally we had to look at the packages and ask ourselves if we could live with what was on the MVS system. In most cases we elected to use the software that the corporate office had already purchased for the Pittsburgh machine," Knapp explains.

In anticipation of having to take on the additional processing load of the 4341, the 4381 was also upgraded to a dual-processor (6MIPS) version. Alas, the results were not good enough.

"We brought it up with MVS under VM, attempted to run two DOS systems on that one 4381 and it kind of fell apart. It was incredibly slow. There was a lot of scrambling to figure out what to do next and it was decided to keep both the processors through the conversion period," Northrup continues.

Another reason to keep both was that the applications staff was, at the same time, getting the corporate databases and systems converted to Cullinet's IDMS database management system. This project, begun before the data center consolidation decision, added significant test load to the 4381 that further hindered the hardware consolidation.

"Another tough decision was to continue (and complete) the database conversion under DOS even though we knew the same applications would have to be converted to MVS later on," Bob Hart, database administrator, points out.

No Sweat Stress Testing

Find out now—not six months from now—how software packages and new releases affect productivity. Use CICS/REPLAY'S unique "Run Unit" design to specify transaction mixes without the need to combine or multiply sessions repeatedly. Determine capacity precisely. *Stress testing made easy.* ▶ CICS/REPLAY lets you review, perfect and execute CICS applications and releases online automatically in just one session. You get a clean cut over into production—and a single person can do it all. ▶ Find out more about how CICS/REPLAY keeps costs down and systems up. Call: Interactive Solutions, Inc. 201-488-3708/212-475-2400

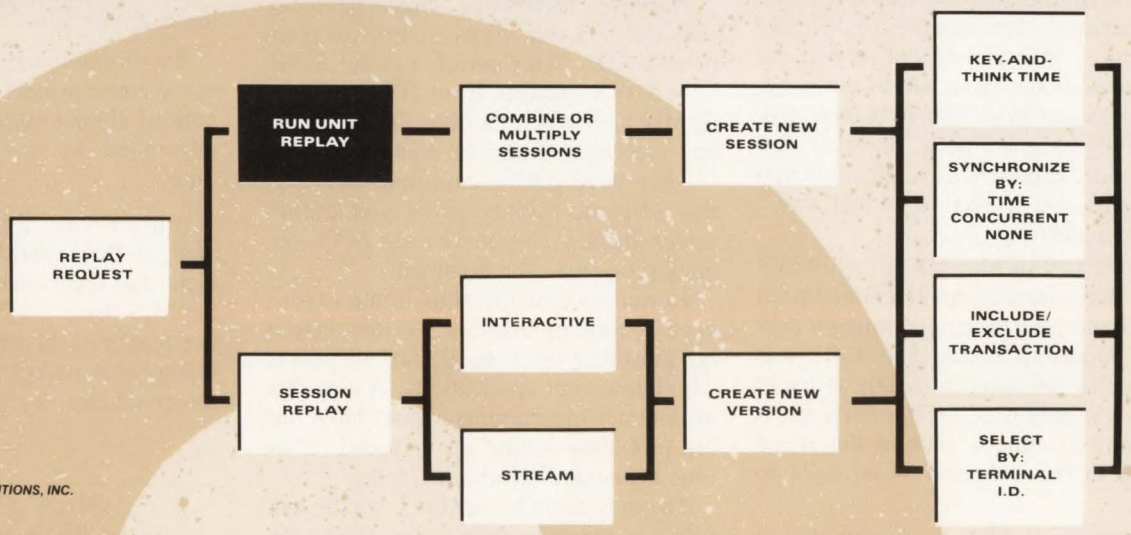

INTERACTIVE SOLUTIONS, INC.
A TBS Company
53 W. Fort Lee Road
Bogota, NJ 07603-1202

CIRCLE #164 on Reader Service Card ▲

An unrelated problem was that Resource Access and Control Facility (RACF) had been used for security on the Pittsburgh 4341. Some applications were secured with RACF and some were not. The people who knew the situation were no longer with the company.

By the end of January, 1988 it was quite clear that the upgraded 4381 was never going to be adequate. So it was decided to bring in a 3081 CPU (a 12MIPS machine) to replace both 43XX boxes. Negotiations with Comdisco, Joy's third-party lessor, resulted in no increase in the monthly lease rate after the swap out.

The 3081 requires water cooling, something Joy's computer room was not equipped for. While the chiller and plumbing were being installed in February, Northrup did the IOCP coding for the 3081 and the input/output device generation. "Luckily, the I/O generation worked the first time," he says.

The 3081 was installed (initially replacing the 4341) by the end of February with all the devices supported. The conversion to IDMS had been completed (under DOS) in November, 1987 and the DOS-to-MVS conversion begun in October under the coordination of Hart and Dave Weaver, senior systems analyst, proceeded through March and April. Parallel testing was carried on and the DOS disk files were made accessible to the MVS system. About 50 applications of varying sizes (with about 1,500 programs) had to be converted.

They chose the big bang over the phase-in approach accomplishing the DOS-to-MVS switch-over in the course of a weekend. "We found that we had too much interdependency between files of different job streams. In a phased conversion the job streams would have had to move between DOS and MVS. There is more risk with the big-bang approach. However after we had more experience testing converted job streams, we decided to go that way," Knapp explains.

"On the first of May we brought MVS and all the converted applications up on the 3081 and we have been running it ever since," Northrup notes. The 4381 was kept around for another month as back-up, but was only used for file retrieval. It was removed June 1 putting the space situation in the computer room back to normal.

The month of May was spent tuning the 3081 system performance. "We have been able to get better on-line response times than we had under DOS even with the consolidated processing load. Eighty-nine percent of all CICS transactions are now complete in less than one second with 94 percent under two seconds.

Looking Back

"Our operating staff was small to begin with and we were telling them that we were going to add a significant number of new applications and remote users and another whole system that they would have to operate without any more people. Despite some obvious concerns about the new workload, they reacted favorably to the opportunity to learn something new and to the challenge of getting this new environment under control," Knapp points out.

Continuous miner and shuttle car.

"A lot of blood, sweat and tears went into this thing but we have saved the company about $800,000 a year by having one instead of two data centers. We now have a single operating system environment and a consolidated communications network to support our expanded client base," Knapp continues.

Northrup, meanwhile, thanks his lucky stars that he came across a particular piece of software. "Our savior was the SDSF 5798-DWX product from IBM, a spool display and search facility. The operators had problems managing queues in MVS — they were getting confused. However this software facility gives you a full-screen display so you can see MVS operating on-line," he comments.

Horton sees two lessons in the experience. "Looking back we see that there is no substitute for a good staff of people who know your applications. A new-hire or contract programmer would have had big problems getting up to speed in the time available," Horton explains.

The only thing brought in from the outside for the conversion was CA-CONVERTER (Computer Associates, Garden City, NY). "By relying on our in-house resources we did not forget what 'outsiders' would have taken away with them at the end of the project. More importantly, the intensive on-the-job learning required to get the job done quickly raised the skills (and confidence) of everyone on the MIS team," Horton points out.

"The other lesson was: do not over-plan. I know that IBM and others have detailed conversion methodologies showing, step-by-step, how to plan and carry out a conversion. If we had followed those models to the letter, it would have extended the project at least six months. We planned as we went along, phase by phase.

"And I think we did a good job communicating with our user clients to let them know the rationale for what we were doing. This is difficult because the work is behind the scenes and there are benefits not apparent. But progress reports kept the project visible and we met with them often to keep them informed," Horton says.

Through it all, the applications development staff under Al Knecht and Tom DePriest also managed to implement a new Cullinet accounts payable system for the division's 11 operating units. Also they implemented a new order entry system using a knowledge-based, expert-system approach to automate the configuring of Joy Technologies' highly optioned mining machinery.

But probably the main accomplishment of the entire project was that the end users noticed almost nothing except for an improvement in on-line system response time.

"In our case this can only be attributed to the individual talents, commitment and teamwork of all the people on the MIS team. In the final analysis, it is the quality of the people and not the technology that makes the difference," Horton concludes.

ABOUT THE AUTHOR

Lamont Wood is a free-lance writer specializing in high technology and mainframe computer-related topics.

THE IBM MID-SIZE MAINFRAME DISASTER RECOVERY SOLUTION

On-Site 3890 Capability

Cary, North Carolina

Tewksbury, Massachusetts

Niles, Ohio

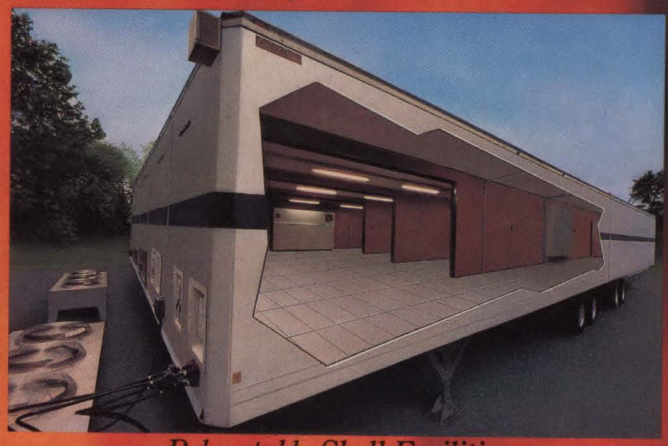

Relocatable Shell Facilities

HOTSITE® offers:

- Multiple, fully equipped recovery centers
- Relocatable shell facilities
- PC based planning services
- Data communications network backup

For more information about HOTSITE® return the coupon or call our Marketing Services Department.

HOTSITE®
A DIVISION OF COMPUSOURCE

HOTSITE® Marketing Services, Cary, NC (919) 460-1234

YES, please send me more information about your disaster recovery services.

Name _____

Title _____

Company _____ *Phone* _____

Address _____

City _____ *State* _____ *Zip* _____

Mail to HOTSITE®, Marketing Services, 110 MacKenan Drive, Cary, NC 27511

CIRCLE #40 on Reader Service Card ▲

MVS Performance Management Legends

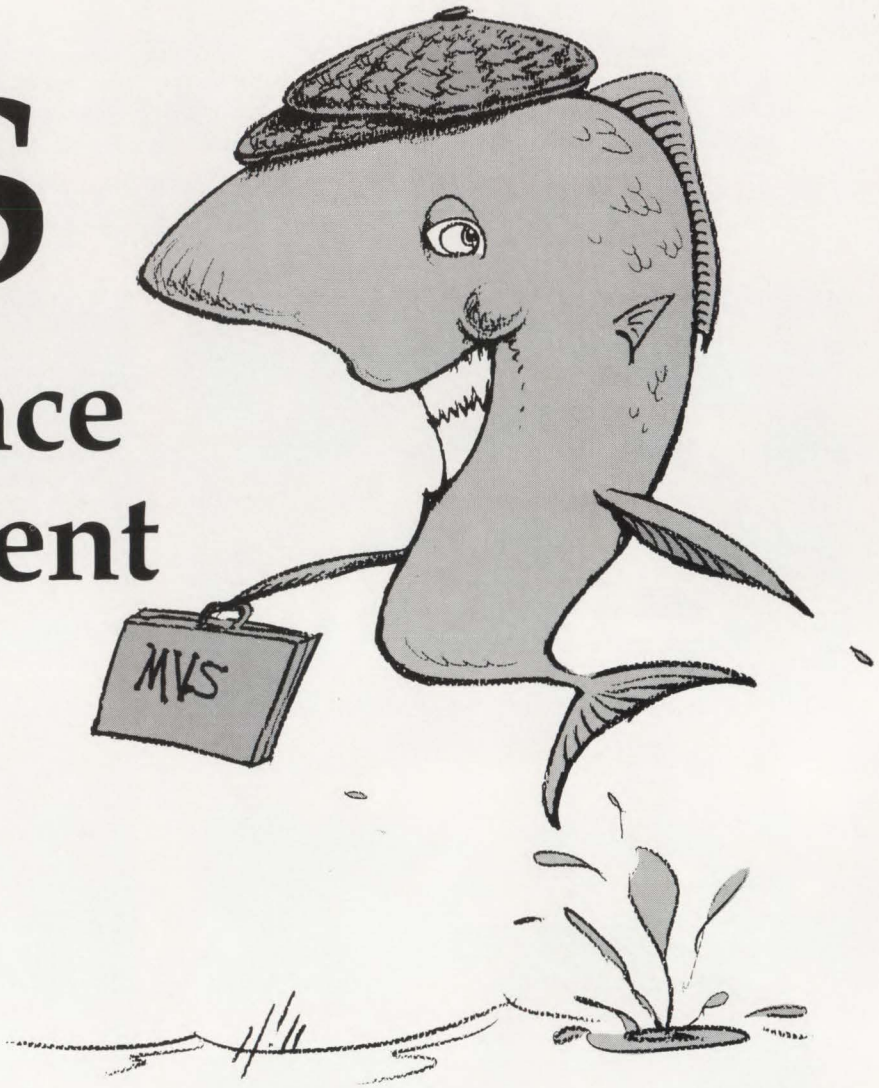

By Stephen L. Samson

The process that creates some of the more popular MVS legends is examined. The intent is to discard the legends that are no longer useful and try to guess which of today's clever observations may become tomorrow's legends.

It is not true that most people dislike reading. Almost anyone will read a tabloid or a magazine that appeals to the nonverbal right side of the brain. We get in trouble when we need to engage the left side of the brain, the part responsible for analysis, deduction and induction. Unfortunately, the way systems programmers seem to be expected to learn about MVS is to read a publication that repels the right brain while overworking the left brain. The *Initialization and Tuning Guide* has no color, no images and illustrations that demand more thought than the printed words. The usual outcome is heavy breathing — but not the kind that goes with physical exertion.

As an alternative, the systems programmer says, "Do not teach me, give me the answers!" The answers being sought are simple imperative statements like, "Do not swim right after eating!" or "Always kick down the accelerator before starting your car!" or "Never run channels more than 30 percent busy!" In a more reflective mood, he or she may welcome declarative statements that sound wise like, "Batch swapping is death" or "Paging is something that just happens to MVS." We receive such input readily and, since its correctness seems to be consistent with experience, the information is reinforced and becomes part of our unquestioned knowledge base.

In systems programming, there is far too much to do: making sure the system is functionally complete with compatible levels across all subsystems; selecting options; writing exits; and then doing it all again when the next release or even the next PUT comes out. There is just not time to curl up with a bad book and arrive at an independent, correct determination of how to define the performance periods of TSO. Who cares?

TIME

"Progressive" (kernel-at-a-time) DOS to MVS conversions are anything but. They tie up your programmers, your system, your productivity for 12 months, 18 months, 24 months — or more.

CTG can do the job for you in only 6 months with CORTEX.

Most companies have converted to MVS in just 6 months with the CORTEX mass migration technology — and you can, too.

COST

The costs of progressive DOS to MVS conversion have a way of "progressing" right out of control. But conversion costs are controlled with CORTEX.

CTG gives you a fixed price with CORTEX. Up front.

You get program conversion. JCL generation. File transfer. Database migration. On-line documentation. On-site service. Follow-up support.

All at a cost you can count on.

RISK

If your "progressive" DOS to MVS conversion turns into progressive confusion, you risk company-wide DP gridlock.

CORTEX mass migration eliminates risks associated with progressive conversion methods. No dual production systems. No freezing of applications development. No lost productivity.

And it works. No "black box." No "vaporware." No "big bang." You get proven conversion technology that we'll demonstrate for you.

For DOS to MVS Conversion in a fixed time, at a fixed cost, with no risk, call

1-800-DOS 2 MVS

CTG/CORTEX

COMPUTER TASK GROUP INC.

3095 Union Road • Orchard Park, New York 14127

(716) 674-9310

CIRCLE #56 on Reader Service Card ▲

As if the sheer difficulty of doing the job were not enough, consider the risk-to-reward aspects of trying to learn what is admittedly error-prone and applying it successfully. If you succeed, no one will notice. If you fail. . . .

It follows, then, that the safest course is to do nothing. However MVS will not let you off that easily. You must do something at least to adapt a preconfigured system offering to the needs of a production workload. If you cannot get away with inaction, it is likely that you will try the next easier approach: attempt to emulate the actions of others who have survived the experience. We are led to choose among the purveyors of prepackaged wisdom and their wares.

The most likely source, IBM, is not likely to supply crisp answers. It would not do for IBM to make a flat statement like, "Do not use swap datasets." This is because a contrived example suggesting the need for swap datasets can be constructed. IBM's prescriptions tend to describe how to do just about anything, but with few value judgments about what is good to do or what is just possibly unwise.

Enter the Out-of-Town Expert

The condensed wisdom of MVS arrives at your shop in the person of the out-of-town expert, also known as Charley the Tuner. Charley has heard all the stories and may have even experienced some of the nightmares. He may even know more than you do. The advantage of such a visitation is that Charley will usually support the perfectly good conclusions of systems programmers or capacity planners who know the environment. He will add or change a few words here and there to sprinkle "expert dust" on what you already know. When Charley speaks, however, managers listen. They have paid him enough so that they must at least appear to accept his recommendations.

Worthy Charleys have read about MVS and have thought about it as well. They analyze your system and look for appropriate solutions in the context of your needs instead of trying to fit your observations into preconceived molds. They are willing, even eager, to share their analysis and justification for each recommendation often to the point of narcotizing the unwary. Other out-of-town experts, not so well qualified, will usually offer advice (often the same advice for all) without explanation, as if sharing one's thought process will diminish the value of the advice or (horrors!) make the expert dispensable. The lack of explanation or justification usually signals lack of knowledge and touches the process of legend formation.

How MVS Legends Are Formed

An MVS legend typically comes about when someone tries to do something new (at least to him or her) in the MVS environment for the first time and discovers that prior analysis and planning might have been a good idea. Instead of humbly accepting correction, the ego insists on justifying the unproductive effort and the incident becomes depicted as an unexpected MVS "behavior." The symptom is magnified, generalized and reported to others in the darkest terms: "MVS cannot control paging!" "TSO LOGON takes forever!" "XA will not fit!"

Some of the first-hand conclusions are true and some reflect lack of understanding. The hysterical broadcasts are invariably oversimplified and colleagues love to revel in them. (Just go to SCIDS at any major SHARE at 11 p.m. on Thursday. The songs that get applause, while they are sung, are the ones that recall days and nights of horror.) As time passes, the hysteria is polished off, the language becomes muted, the number of people the word has spread to rises (as in a pandemic) and rumor becomes legend.

Some Legends Examined

"MVS Uses Storage Inefficiently"

The Germ of Truth

When the first release of MVS came out, the systems of the time barely had enough storage to IPL. Single Virtual Storage (SVS) had run well on systems of 768K or 1MB. MVS had much larger fixed data structures to accommodate and thus a larger initial real storage requirement. MVS needed up to 2MB to get started and significant workloads required four or more megabytes. What might have happened was that the doubling of real storage requirements was not exactly palatable to the SVS installations being persuaded to jump on the MVS bandwagon and the truth came out slowly, incrementally and looking like rumor. When IBM finally made realistic real storage requirements for MVS known, the buildup of skepticism took a while to peel off.

Later Developments and Second Thoughts

MVS' appetite for real storage continued to grow. However as its use for application code began to far outweigh its use for system code, MVS systems pro-

VSAM-LITE
Automatic Data Compression for VSE
- Saves VSAM space by 50% or more
- Reduces DASD I/O activity...improves system performance
- Improves on-line response times
- Easy to use...requires no JCL or program changes
- Analysis utility forecasts actual space savings

Call now for your **FREE** copy of the Compression Analysis program that will show how much disk space your installation can save by using VSAM-LITE.

CALL...

Toll free **(800) 243-9542**
in Conn. **(203) 792-5100**

Universal Software inc.
Brookfield Office Park
Brookfield, CT. 06804

grammers came to recognize that MVS' use of real storage is efficient but many applications are extravagant and wasteful. Early MVS subsystems like VSPC attempted elaborate schemes for private real storage management. When they were written (and *where* they were written), MVS' services and capabilities were not known and such defensive measures may have been necessary. Over time, as subsystem and application developers got to understand MVS and the essential elements of coding for efficient use of real storage in a virtual storage environment, the ''problem'' receded.

Lesson of the Legend

What can we learn from this legend?

- New systems require new expectations
- The truth takes longer to sink in when it is held back at first
- Early subsystems of a new system may not support it well at first
- Systems get well before subsystems
- The sins of the subsystems are visited upon the systems.

"TSO and DB/DC Do Not Mix"

The Germ of Truth

In early MVS (before the introduction of swap datasets and, of course, before logical swapping) the TSO swap load so dominated the paging devices and paths that paging response times for other workloads could not be expected to be fast or consistent. Contention for the CPU on uniprocessors was a problem as well, especially when the person supporting each subsystem tried to make his or her charge look as good as possible. The TSO paging activity caused inescapable high-priority overhead, regardless of TSO's dispatching priority.

Later Developments and Second Thoughts

The early reaction was simply to dedicate systems to either TSO and batch or to CICS, IMS DB/DC or similar transaction processing subsystems. Another early response was to try to fix the symptoms. The use of fast paging devices (''drums'' and later their solid-state replacements) was a way to sustain high paging rates with good response times. ''Directed paging'' modifications to MVS attempted to segregate the paging activities of major workloads.

As systems grew larger, it became necessary to reconsider the viability of mixed workloads. The successive changes in MVS' swapping algorithms (swap datasets, logical swapping, extended swap, contiguous slot allocation, storage isolation and finally the use of expanded storage) successively unloaded local page datasets and facilitated the growth of MVS systems with noninterfering mixed workloads. Growth in real storage sizes moved along with the trend often leading, sometimes lagging, but finally emerging in the lead. Multiprocessors, as well as MVS/SE2 SRM improvements, eased CPU contention problems (the latter by making possible centralized control of dispatching priorities) and XA's outboard channel subsystem facilitated offloading of high priority overhead.

Lesson of the Legend

The MVS real storage experience has conditioned expectations in many areas.

EASY ACCESS TO REPORTS IN THE SPOOL QUEUES with the JES2 QUEUE STRIPPER™ or POWER QUEUE STRIPPER/LOADER™

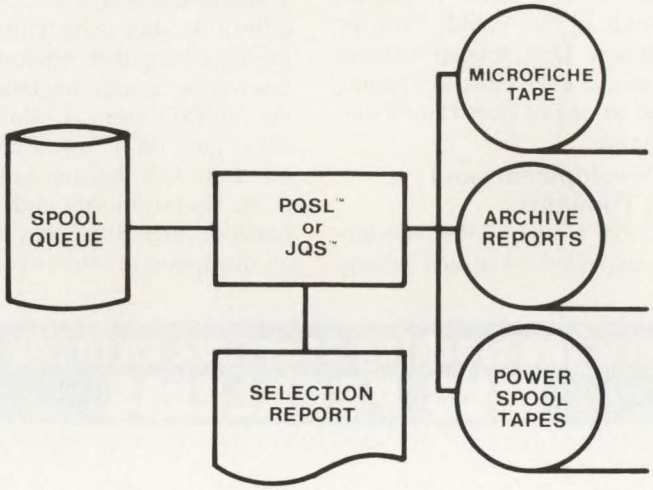

- Supports MVS/SP, MVS/XA, DOS/VSE, VSE/SP
- Retrieves reports directly from JES2 or POWER spool queues.
- Alters report characteristics after processing
- Multiple report selection criteria
- User Exit facility provides easy access to reports

Find out more about the different uses and features of JQS™ and PQSL™. Call or write today!

SERCON corporation

4611 North Lindhurst Dallas, Texas 75229
(214) 691-3020

CIRCLE #77 on Reader Service Card ▲

Storage sizes in personal computers, for instance, have moved ahead with great rapidity as OS/2 moves into view. That operating system should profit from the MVS experience. Ready solutions should be at hand when the dependency on a single paging device causes OS/2 multitasking workloads to interfere with each other due to paging contention.

The CPU contention aspect of accommodating mixed workloads seems to be settled with the offload of I/O processing CPU activity in XA and ESA along with large MP configurations. CPU contention will certainly reappear in UP systems from the 3090-180E down to superminis and micros. The search for effective and efficient dispatching schemes will continue.

"MVS Needs Fast Paging Devices"

The Germ of Truth

The history discussed above identified a past period when the real storage available to MVS was insufficient to accommodate a workload that could drive the CPU to saturation. One stopgap measure was to try to mitigate the effects of paging by making the paging devices fast for single-page retrieval.

Later Developments and Second Thoughts

Since MVS/SP 1.3.0 and the introduction of both extended swap and contiguous slot allocation, the need for fast paging devices in MVS has diminished to disappear with the acceptance of expanded storage.

Lesson of the Legend

Whenever real storage is insufficient to support the workload of a fully utilized processor complex without noticeable paging delay, paging problems will reappear. As the incidence of such problems diminishes, the performance impact for any such episode will be more keenly felt. Under these conditions, the search for solutions may bring back the demand for fast external paging devices. Beware and prevent!

"CPU Hogs Hurt Other Work"

The Germ of Truth

Back in the bad old days before the ICS and APGRNG, those submitting batch jobs could specify their dispatching priorities. A scientist ignorant of the natural order of things in MVS might try to run a Monte Carlo analysis through 100,000 rolls of the dice at high dispatching priority, reasoning that whatever the director wants *Now!* is most important. Never mind the 50 TSO users (I said it was the old days) now on a forced break and never mind the fact that the same job running at the lowest priority in the system might finish in only 10 percent more time with no disruption to other work.

Later Developments and Second Thoughts

We have learned to control the relative dispatching priorities of all work in the system and we know how to use mean-time-to-wait groups to ensure optimum throughput. We may not have learned how to calm down the hysterical batch submitter, but we are working on that.

Lesson of the Legend

This is a proposition that can be argued from either position. Stated as above — a pure number cruncher with good real storage behavior — it is absolutely false. On the other hand, if the CPU hog likes to page or do I/O, it is a different matter. The system crippler might be a BSAM (without chained scheduling and 100 READs before CHECK) concordance program running against the King James Bible on an otherwise lightly loaded CPU with its dataset blocked 80-80 (frequent I/O) on a solid state DASD (fast I/O with no disconnect) and incorrectly specified storage isolation. Anything else in the system will go to sleep even if the professor's program is at M0.

That all-purpose answer, "It all depends," applies well here.

"Swapping Requires Swap Datasets"

The Germ of Truth

When swap datasets were introduced in MVS, their purpose was to offload LSQA page movement from the locals. For the cost of a single swap dataset, 50-100 pages per second might be redirected. Swapping improved so much through this change that the use of a swap dataset became standard in TSO systems although never an absolute requirement.

Later Developments and Second Thoughts

The improvements in MVS physical swapping, cited previously, outdistanced the capabilities of swap datasets. Extended swap, made available first in an IUP for SP 1.1 and incorporated into SP 1.3, redefined swapping as a single stage operation. The entire swap group moved together to swap datasets if present or to local page datasets without swaps or if the swap capacity was exhausted. When swapping to locals, contiguous slot allocation (another SP 1.3 innovation) ensured the creation of channel programs comparable in efficiency to those used for swap datasets. Subsequent enhancements

are you working overtime on these MVS PROBLEMS

- IDCAMS DID NOT DELETE A CATALOG ENTRY PREVENTING YOU FROM REDEFINING AND RESTORING A FILE.
- VSAM FILES EXIST BUT CANNOT BE SCRATCHED BECAUSE THEIR CATALOG ENTRIES ARE MISSING.
- A CATALOG MUST BE RECOVERED
- VSAM FILES ARE OVERALLOCATED AND YOU NEED THE UNUSED VSAM DASD SPACE **NOW.**
- A PDS MEMBER HAS INADVERTENTLY BEEN DELETED OR A PREVIOUS VERSION OF AN UPDATED PDS MEMBER MUST BE RECOVERED.
- VSAM FILE INTEGRITY ERRORS HAVE YOUR ONLINE SYSTEM OFFLINE.
- YOU GUESS AT THE CI SIZE AND THE AMOUNT OF DASD SPACE TO ALLOCATE WHEN DEFINING A VSAM FILE.
- POORLY BLOCKED FILES ARE USING UP YOUR CPU CYCLES
- A VTOC MUST BE EXPANDED OR RELOCATED IN THE NEXT 60 SECONDS

DP TECHNICIAN EASILY RESOLVES THESE PROBLEMS AND PROVIDES FACILITIES FOR:

- DASD MANAGEMENT
- DISASTER/RECOVERY
- VSAM MANAGEMENT
- CATALOG MANAGEMENT
- DATA SET MAINTENANCE
- DASD MIGRATIONS/CONVERSIONS

dpt

FOR MORE INFORMATION CALL OR WRITE:
DATA PROCESSING TECHINQUES, INC.
21 LAKEVIEW DRIVE
WEST ORANGE, N.J. 07052
201-325-3251

CIRCLE #98 on Reader Service Card ▲

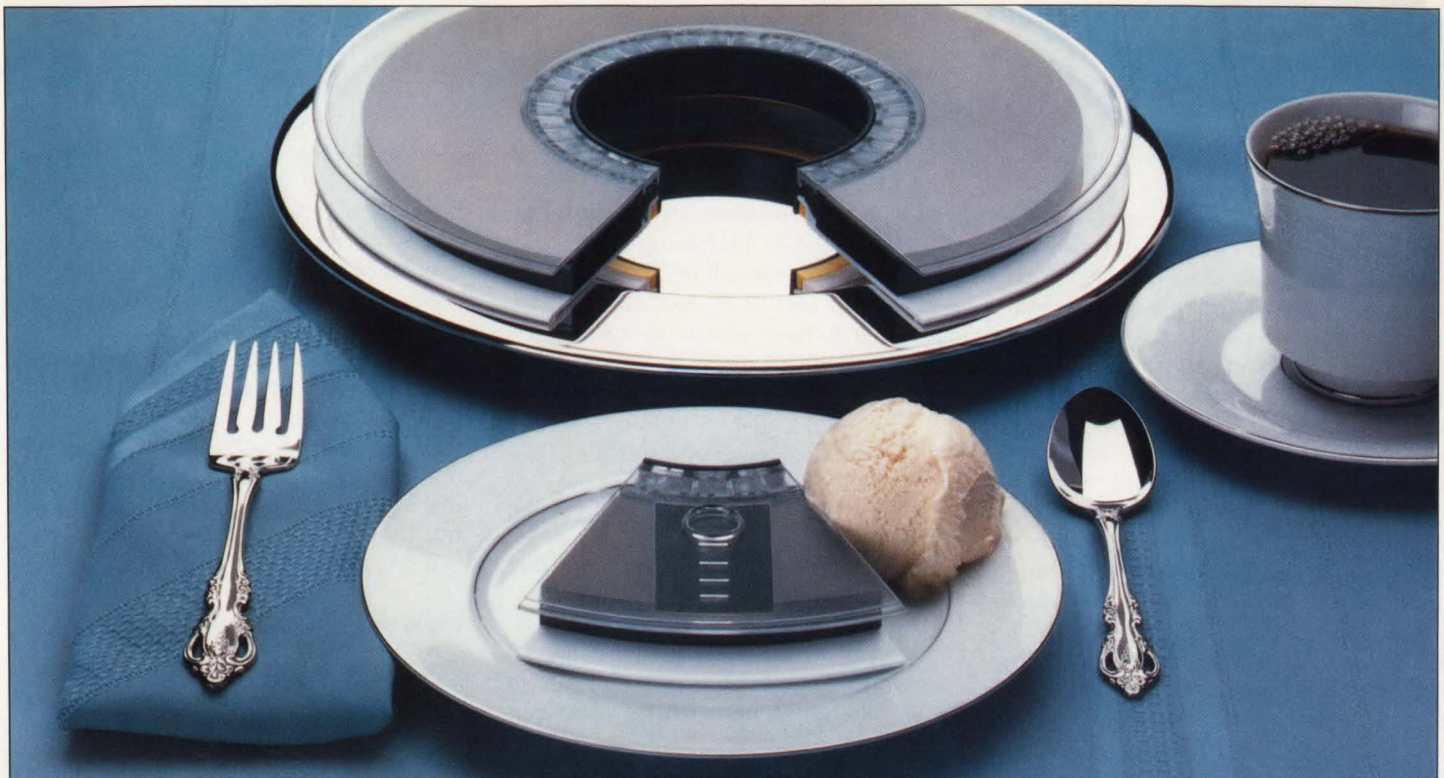

When choosing a network access package...why settle for just a piece of the pie?

If all you want is access to your network, a number of software vendors can help you. But, if you want more than just a piece of a product, you need NCI from Westinghouse.

Network Control Interface gives you the power of total network control. It is the only network access and control product available for MVS, MVS/XA, VSE and VM operating systems. In fact, regardless of the size and complexity of your VTAM network, NCI provides all your users with a single system image, with identical entry screens and procedures.

Of course we give you single point entry for quick application access at the push of a single key. That's the easy part. But we also provide application load balancing, logon queueing, unlimited exit processing and a dialog manager...powerful features to improve your productivity...without sacrificing system performance.

NCI can read the contents of your security data base and dynamically build user menus, listing only applications the user is authorized to access...reducing user confusion and simplifying security administration.

With NCI, you can customize menus to meet your specific user needs, or use some of the over 100 sample panels provided in our Starter System...the options are unlimited.

To truly appreciate the power of NCI, we invite you to call or write us for a free, 30-day trial.

More than just simple network access...NCI from Westinghouse.

Westinghouse
SOFTWARE **SOLUTIONS**

Westinghouse

Management Systems Software
P.O. Box 2728
Pittsburgh, PA 15230
(800) 348-3523
(412) 256-2900 in PA

CIRCLE #128 on Reader Service Card ▲

corrected early inefficiencies in the algorithm.

The increased number of pages moving to swap datasets in conjunction with the fixed swap set size of 12 slots meant that the number of swap sets per swap would increase from one to four, five or more. To ensure best possible swapping performance would require many swap datasets on separate paths, as many as there were swap sets in an average swap group. As TSO working sets grew, the number of needed swap datasets continued to grow. Finally, some Charleys started recommending allocating only half the number of swap datasets needed for full concurrency. Because the data rates became high, these needed to be on dedicated devices. A "technopolitical" question arose because swap datasets are small, rarely more than two hundred 3380 cylinders, yet the device had to be dedicated for acceptable performance. How could the systems programmers appropriate from four to as many as ten 3380s for swapping without accounting for the unused space?

However local page datasets on dedicated devices are more efficient for swapping than swap datasets, especially if the devices are 3380s. Increasing the number of locals while spreading the swap load across them increases operational flexibility. When the workload swings from a swap-dominated character to one in which paging is more prominent, the paging configuration adapts smoothly to the change. Swap datasets would stand idle while paging devices worked harder.

IBM has finally come out with a flat recommendation against swap datasets — on systems with expanded storage. In that environment, they are virtually unused. IBM experts, including Tom Beretvas, continue to say that there are some special environments in which swap datasets have some value, but modeled or measured proof is not offered.

Lesson of the Legend

The legend is dead wrong. MVS systems never needed swap datasets in the sense that they need page datasets. A "normal" MVS configuration has no swap datasets. When severe storage constraint appears on a system with a rough balance between swappable and nonswappable work and the nonswappable subsystems cannot be protected from page faulting adequately with storage isolation, then and only then should swap datasets be considered. Even then, since additional paging resources will be needed anyway, more locals should be tried first.

When all else fails, add swap datasets (enough for full concurrency to minimize idle storage occupancy) on separate paths making sure to specify generous storage isolation for first and second period TSO.

The whole question may usually be avoided by increasing real storage.

"Batch Swapping Is Trouble"

The Germ of Truth

At the same time that legends about TSO swapping were being formed, batch swapping was being portrayed as even more evil. The root cause was the same: the I/O load caused by swapping swamped the page datasets causing poor response time for demand paging and for more urgent (TSO) swapping. Batch swapping hurt TSO response time, not to mention the response time of CICS or IMS on the same system.

Later Developments and Second Thoughts

The original design of the SRM used exchange swapping as a means of giving competing batch jobs an equal chance of running to completion when storage constraint kept the multiprogramming level (MPL) less than demand. Once batch swapping had acquired its unsavory reputation, heavy-handed measures were taken by various groups within IBM to make exchange swapping virtually impossible. (The measures included a default ISV of 100,000 service units and a workload scale in the IPO IPS running from one to 10, rather than one to 100, ensuring a low probability of a Recommendation Value (RV) difference greater than one — the criterion for an exchange swap.)

Those with heavy hands did not consult the Auxiliary Storage Manager developers. Just as exchange swapping had been virtually prohibited, the SP 1.3 rewrite of the ASM made batch swapping quite respectable. Later developments such as XA, the SP 2.1.2 ASM enhancements and the introduction of expanded storage have

SYSTEM PERFORMANCE

Let members of our full time professional staff help you improve performance of your systems.

-MVS -VM -VSE -VTAM

-CICS -IMS -VSAM -COPICS

Planning for and Implementation of the above Products, Hardware Evaluations, Relocations, Consolidations and more. Call Today.

DOWLING ASSOCIATES, INC.

Information Systems Specialists

OHIO	CONN.	NEW YORK	FLORIDA
(216)	(203)	(212)	(305)
656-1136	255-6316	912-1212	782-9522

CIRCLE #83 on Reader Service Card ▲

made batch swapping an ordinary activity in MVS, not a disaster.

Lesson of the Legend

The lesson of this legend is that nothing in MVS performance stays good or bad forever. Fixes designed as short-term become institutionalized and preclude flexibility in the future because the options they rule out are rarely re-examined.

" 'Out-and-Ready' Is Unconditionally Evil"

The Germ of Truth

The RMF CPU report shows an overall number for "Out-and-Ready" (RDY) address spaces. Many analysts have fastened on this number as a key indicator of system tuning effectiveness and have recommended a rule-of-thumb maximum such as "not greater than 1.0 per active processor." If all of these RDY address spaces were TSO users, this guideline might make sense. However the number reported by RMF is system-wide and it does not differentiate between batch and TSO RDY.

Later Developments and Second Thoughts

Since the rule of thumb was advanced, we have come to recognize that we can control individual parts of real storage occupancy. With tools like the Cross-Performance-Group (XPG) report in Candle's DEXAN/MVS or EPILOG/MVS, we can see how the RDY state is distributed across workloads. Then we can adjust domain constraints and contention index algorithms to ensure that the appropriate work is relegated to the outqueue. Once the impact of RDY on key workloads (seen as "SRM (MPL) Delay" in EPILOG or DEXAN detail displays) is at an acceptably low level, the amount of delay to less critical workloads is irrelevant. In fact, in a system with a significant batch workload, the *absence* of some level of RDY might indicate inappropriate MPL controls.

Lesson of the Legend

Tracking the RDY number over time provides a rough index of real storage loading. A rapid increase may signal impending saturation and the need for an upgrade. The overall RDY number is useless for tuning assessment because of its aggregate nature. As in many other areas, a useful capacity planning metric is inappropriate for day-to-day performance management.

"'Overinitiation' Is as Bad as It Sounds"

The Germ of Truth

When words were invented to describe common MVS conditions, the word chosen to mean "the batch MPL is less than the number of open initiators with ready jobs" was *overinitiation*, suggesting the number of initiators is excessive. This pejorative term leads operators to drain initiators lest the performance analyst begin making accusations of *misprision of overinitiation* or perhaps *suborning overinitiation*.

Later Developments and Second Thoughts

Suppose we redefine *overinitiation* to mean "keeping standby batch on the page

Advanced 3480 Technology . . .

"Picture This"

- 1/2 the cost
- 1/2 the space
- Twice the Reliability

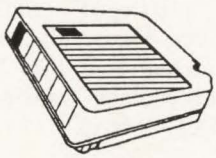

The IBM® compatible cartridge product designed for your budget. At last, an affordable Square Tape Solution, The FAST A-480!

For Information Call: 1-800-527-3237
In California, Call: 1-619-487-8030

fast™ First Alliance Software and Technologies, Inc.
11770 Bernardo Plaza Court, San Diego, California 92128

IBM is a registered trademark of International Business Machines

CIRCLE #107 on Reader Service Card ▲

datasets to increase throughput by taking advantage of irregular demand patterns.'' As batch swapping became less fearsome and performance specialists became proficient at controlling MPL by monitoring the page fault rate (RCCPTRT in the OPT), the swapped-out batch became the principal load balancer in the system. Overinitiation is *required* to make the control mechanism work. To steal a good IBM quote, ''Know who your loved ones are — and always have someone else to kick around!'' (Credit to Ray Wicks and/or Siebo Friesenborg of the Washington Systems Center. I do not know who said it first.) Swapped-out or marginally swapped-in batch is the guy wearing the ''Kick Me'' sign.

Lesson of the Legend

Do not let a bad-sounding word get in the way of a good idea.

''Paging Is Something That Just Happens to MVS''

The Germ of Truth

MVS with default SRM controls will page as soon as real storage becomes fully utilized. When the available frame queue is depleted, real frame replacement (known to the rest of us as page stealing) begins. Subsequent re-reference to stolen pages causes page faults and page-ins. Also, TSO sessions will be logically swapped by default, increasing the demand for real storage and driving paging rates up. Unless the performance specialist gains control of the mechanisms controlling real storage assignment and reassignment, paging will indeed ''just happen'' to MVS.

Later Developments and Second Thoughts

As analysts and systems programmers came to understand MVS' control mechanisms, they learned to keep MVS from exhibiting its default behavior and instead to control MVS paging. Control of MPL based on page fault rate, correct use of storage isolation to regulate swap trim as well as demand paging in subsystems and (in systems without expanded storage) adjusting logical swapping to balance TSO response time against paging elsewhere are the means of gaining control of MVS' real storage management and suppressing its default behavior.

Lesson of the Legend

MVS default behavior will reappear as soon as control limits are no longer effective. Paging ''just happens'' when it is not controlled.

''Logical Swapping Is Good — Or Bad''

The Germ of Truth

The introduction of logical swapping in MVS/SE 1 came at a time when the availability of real storage was, for a time, ahead of demand and IBM could devote its enhancement efforts to improving MVS CPU utilization. Logical swapping used excess real storage to defer physical swap-outs for TSO address spaces, improving TSO response time with few negative consequences.

Later Developments and Second Thoughts

As MVS moved through fat and lean real storage periods, logical swapping waxed and waned in popularity. The default SRM controls favored logical swapping heavily, so when storage constraint appeared, logical swapping became a villain with idle TSO sessions holding frames while crucial active workloads could have used them to reduce paging delay.

With MVS/SP 1.3 and again in early XA, severe storage constraint led to recommendations to cut back severely on logical swapping. Finally with expanded storage, logical swapping became subsumed under the term *processor storage swapping* and is now seen as very good indeed.

Lesson of the Legend

IBM should have learned not to set defaults based on a transitory presence or absence of constraints. We all should have learned that logical swapping is a neutral mechanism that does what it is designed to do under appropriate controls and that the controls need to be set to reflect one's own needs, constraints and priorities.

''Total Paging Rate Is an Indicator of Storage Constraint''

The Germ of Truth

When all paging was of the same kind taking place on the same devices, it was useful to examine the total paging rate on the local page datasets to get an index of storage constraint. This was strictly true when the legend, ''Paging Is Something That Just Happens to MVS,'' was in its infancy — no logical swapping, no controls tying MPL and page fault rate together.

Later Developments and Second Thoughts

Paging is no longer homogeneous and no longer an unconstrained consequence of workload interacting with the real storage resource. Because paging is a controlled variable, no significance can be attached to any paging rate, especially when dissimilar components such as swap working sets, swap trim, VIO windows and page steals are in the mix. With expanded storage, the picture is more muddled and the books do not balance. With MVS/ESA and the introduction of data spaces and HiPerSpaces, ''paging rate'' may cease to be a meaningful term.

Lesson of the Legend

We cannot yearn for simplicity when reality is complex. Note also that the Available Frame Count (AFC) is controlled as well. Therefore, that AFC is equally unsuitable as a measure of storage availability.

''IBM's Defaults Must Be OK''

The Germ of Truth

IBM's developers had a rationale for each default in the members of PARMLIB that control the SRM. In each case, the choice makes sense in the limited context of its selection. Thus, the division of IBM that makes CPUs as well as MVS sets a limit on the amount of CPU resource devoted to ''demand paging'' — no more than five percent. This limit translates to the PAGERT1 and PAGERT2 numbers in the default OPT. Unfortunately, this number (more than a thousand pages per second for large 3090s) has no relationship to the paging rate that the rest of the system (that is the page datasets) can sustain.

As another example, the default dispatching priority is M0, the lowest in the system. Another default is that PVLDP (the dispatching priority of initiators, LOGON and other ''privileged'' address spaces) is unspecified. The result is that PVLDP is M0 by default and TSO LOGONs and batch job and step transitions can die when the CPU is busy.

Later Developments and Second Thoughts

Every MVS systems programmer and performance analyst comes to realize that IBM's defaults for controlling the performance of MVS have only one virtue — they are syntactically correct. From the earliest days of hip-pocket ZAPpable ''happy values,'' every out-of-town expert has had his favorite OK numbers.

FIGURE 1

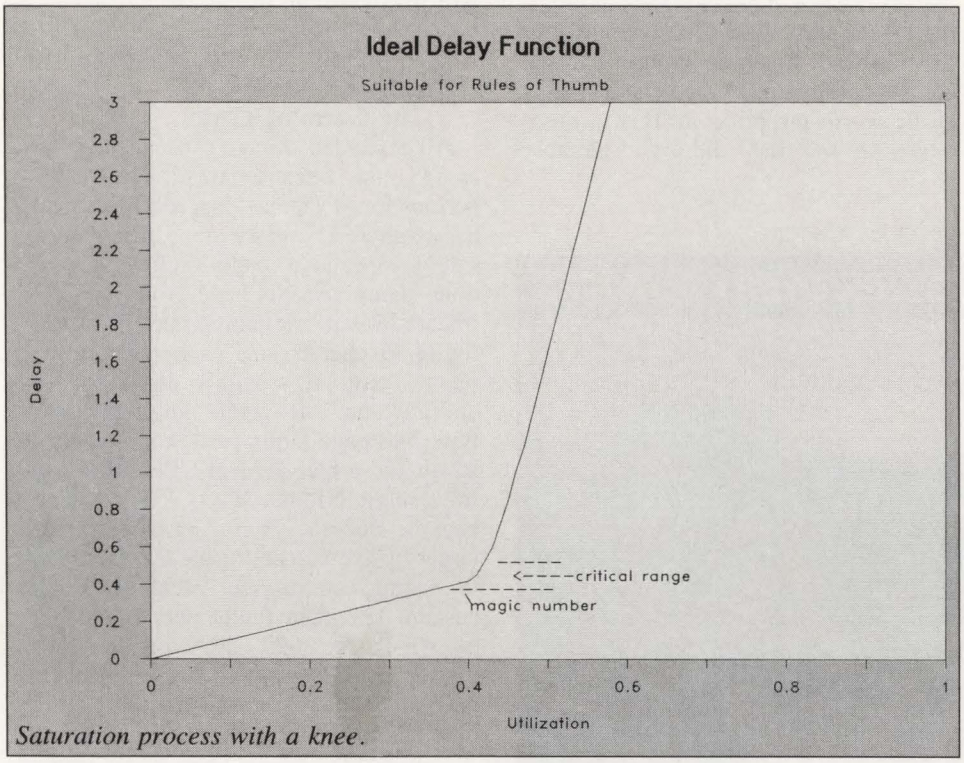

Saturation process with a knee.

Lesson of the Legend

Throw this one out! IBM's defaults cannot help and may hurt. In this area more than any, a book, a class, a consulting visit or (shudder) some time spent studying and analyzing your workload can pay off with a system operating under control within limits you establish.

"MVS/XA Uses Storage Inefficiently"

The Germ of Truth

Here we go again! In an exact repeat of history, MVS/XA's larger fixed data structures required a large increment of real storage over that needed for MVS/SP 1.3. This time IBM attempted to warn MVS customers and held MVS/XA back for a while as a special system for only the largest MVS users.

Later Developments and Second Thoughts

The transition from MVS/SP 1.3 to MVS/XA was so simple that many MVS shops went to XA without a hardware upgrade. They struggled for a while until

TEN GOOD REASONS...

to Let ACTS Corporation Handle Your Technical Training

1. Technical Awareness for Managers Series.
4 one day seminars: Systems Basics, MVS/XA, SRM,SNA/VTAM.
4 days/$250 per day Sept 12 Hilton Head, SC Oct 11 Washington
(Sponsored by Duquesne Systems, Inc.)

2. MVS Overview.
2 days/$500 Sept. 19 Dallas Oct. 6 Washington

3. MVS Systems Programming.
3 days/$750 Sept. 21 Dallas Oct. 12 Washington

4. MVS Architecture (Structure & Logic).
3 days/$750 Oct 3 Dallas Oct 24 Washington Dec 13 Los Angeles

5. Exploiting the MVS/XA SRM Controls.
3 days/$750 Sept 21 Dallas Oct 31 Washington

6. MVS Internals and Diagnostics.
4 days/$1000 Dec 5 Dallas

7. VTAM Concepts & Facilities.
2 days/$500 Oct 6 Dallas Oct 27 WDC

8. Capacity Planning.
3 days/$750 Oct 3 Washington

9. SAS® Basics Plus.
3 days/$500 Sept 26 St. Louis Oct 24 Washington
Nov 14 Portland

10. Automated Operations: Issues and Answers.
1 day/$295 Sept 16 Hilton Head Sept 21 San Francisco Sept 22 LA
Oct 18 Toronto Oct 19 Chicago Nov 10 New York Nov 29 Honolulu

"It's our instructors who make the difference. Once you attend one of our seminars, you will become a regular customer." Bill Carico
PRESIDENT ACTS CORPORATION

® SAS is a registered trademark of SAS Institute, Cary, N.C.

All MVS courses are current for MVS/XA version 2.2, All VTAM courses are current for ACF/VTAM version 3.1.1

ACTS CORPORATION

FOR REGISTRATION CALL
(800) 950-ACTS
FOR OTHER BUSINESS CALL
(512) 693-7900

P.O. BOX 180
KINGSLAND, TEXAS
78639

CIRCLE #123 on Reader Service Card ▲

they had adequate real storage to meet the base requirement of XA. After XA was announced, IBM announced and shipped larger real storage sizes for the 308X processors. In all, IBM's announcements and support prepared users for XA far better than they had been prepared for MVS.

Lesson of the Legend

Again, we are reminded that support of larger configurations has its cost in higher fixed overhead. Enabling applications that were not possible in MVS Version 1 may well be worth the price. In this instance we also see that IBM did a considerably better job of preparation and initial support than they did in the early '70s.

"Expanded Storage Causes High CPU Overhead"

The Germ of Truth

All expanded storage activity shows up as CPU use. Because IBM (as in spokespersons for MVS) has doggedly stuck with the original 75 microsecond figure for a safely repeatable page fault resolution time, many analysts have projected CPU impact for various paging rates using that figure to characterize each page movement. Assuming roughly equal page-in and page-out rates while ignoring migration, 500 page faults per second result in seven-and-a-half percent CPU utilization for a single original 3090 CPU. If we accept the Default System Definition (DSD) target of five percent of the CPU devoted to paging, an original 3090-180 could sustain 333 page faults per second and each E-model 3090 CPU from the 180E and up could sustain 358 pages per second (at 70 microseconds per page) based on cycle time ratios.

Later Developments and Second Thoughts

More information has come out about expanded storage since the original announcements. At a recent user group meeting, for instance, DSD's President James Cannavino spoke of the bus between expanded storage and the CPU as capable of transferring 220MB per second. From this we can infer a burst-mode page transfer time of 17.8 microseconds.

Since many transfers to expanded storage and some from it occur in groups of several pages, the effect of grouping will show up as a time much closer to 18 microseconds than to 70 for subsequent transfers in a group. Let us assume an effective group size of 1.3 pages; in this case the average transfer time is reduced to about 58 microseconds and the five percent CPU point is reached at 432 page faults per second per CPU.

Lesson of the Legend

In the case of expanded storage measurements, the truth will come out and it is probably better than the original story. It is likely that IBM was careful to condition expectations for expanded storage to guard against the possibility of early negative experiences. The new role of expanded storage in Enterprise Systems Architecture (ESA) seems to justify the cau-

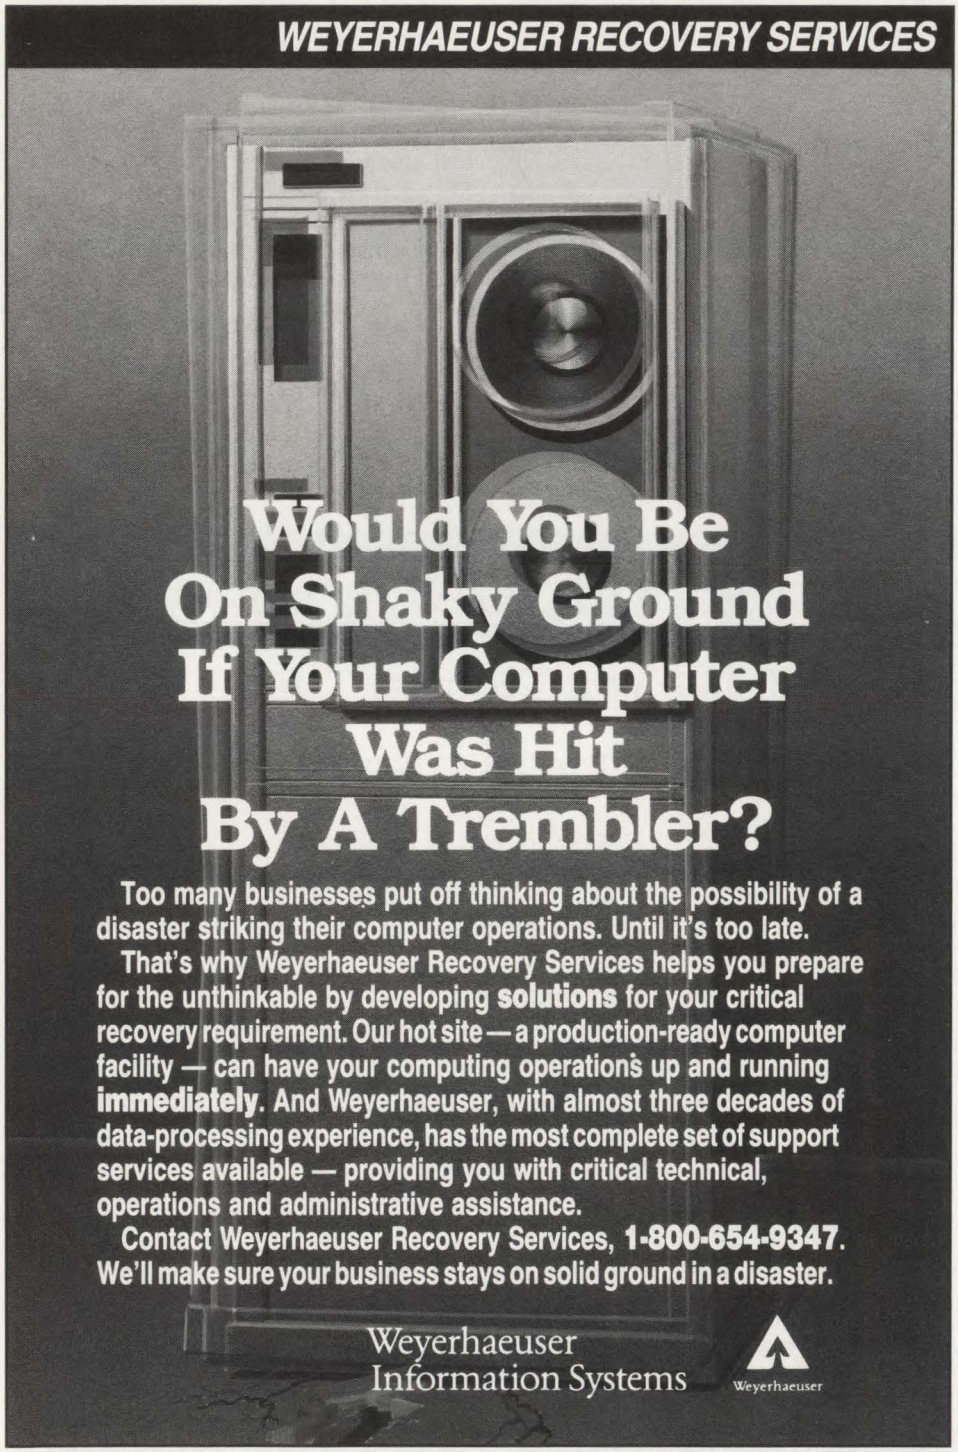

FIGURE 2

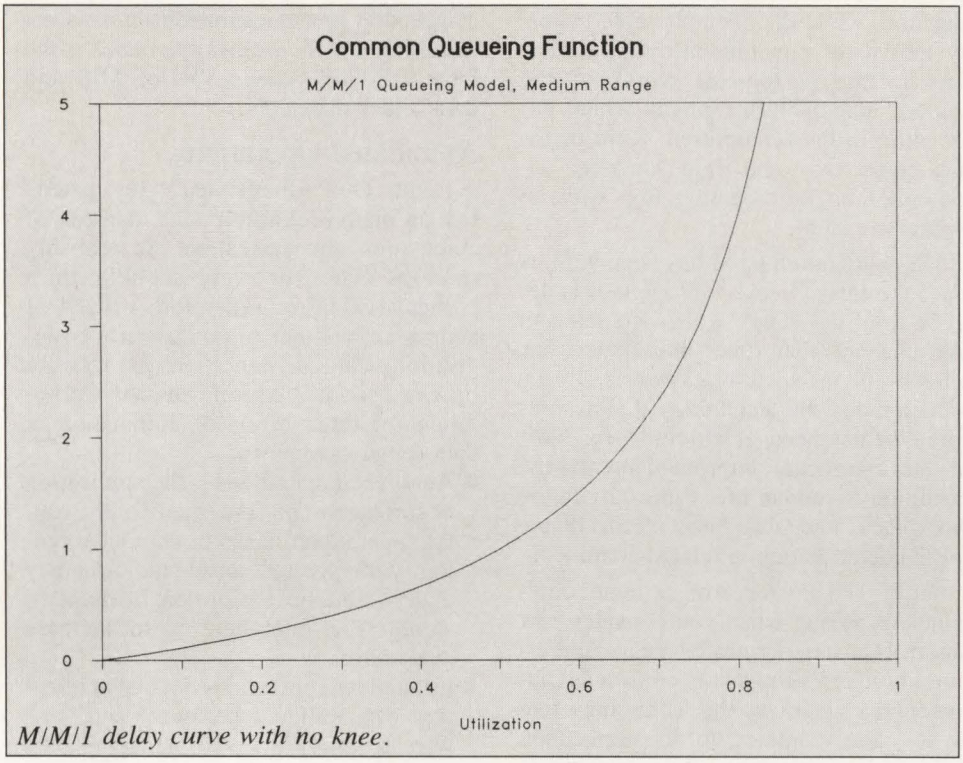

M/M/1 delay curve with no knee.

FIGURE 3

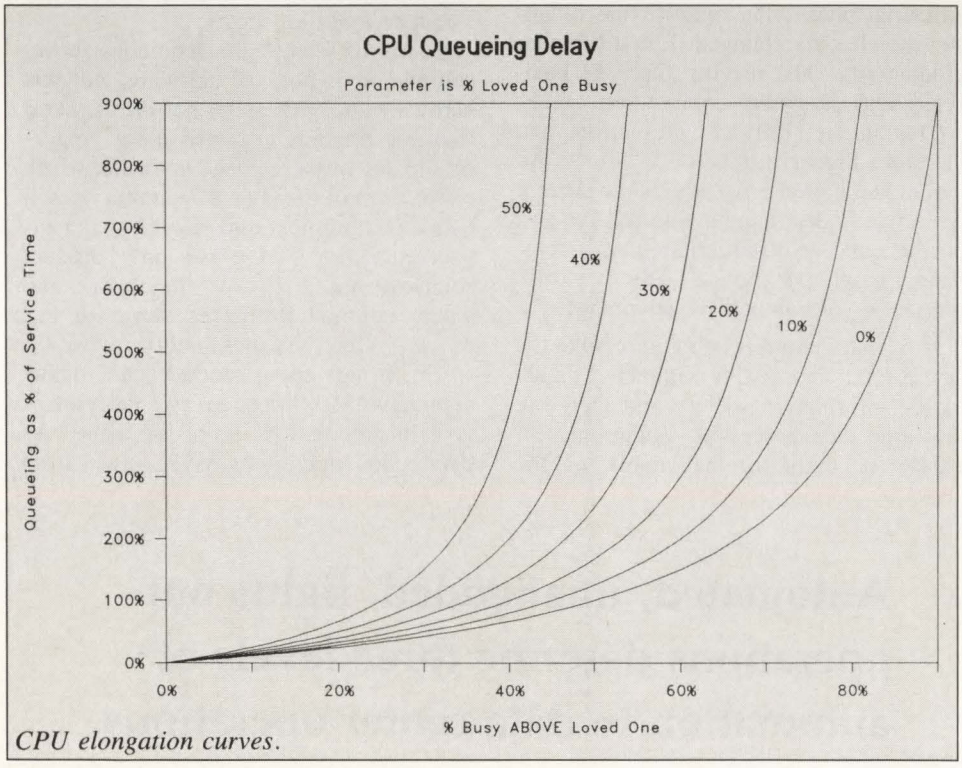

CPU elongation curves.

tion. In the time-honored manner of legend formation, the lack of information led to suspicion, rumor and a new legend. (See the last legend in this paper.)

"VM/XA Can Support Production and Test MVS/XA"

The Germ of Truth

From its beginning, VM/XA System Facility has been able to support the operation of one or more MVS/XA virtual machines. Unfortunately, they have been supported with a VM bigot's understanding of MVS. "Certainly you can have a V=R system with SIE assist for efficient I/O. We call those dedicated devices." "Oh, you want to share those devices with another system? Certainly!" "With another virtual machine? Oh, no — 'dedicated' means what it says: dedicated to a single virtual machine and not accessible to any other. Of course, you can share across virtual machines if the devices are CP-owned. . . ."

Later Developments and Second Thoughts

VM/XA has not been and probably will never be able to support a typical MVS need: a production system sharing DASD with a test system, all devices operating with native efficiency. To do so would require accommodation from VM to the needs of MVS and this is highly unlikely.

Fortunately, a more appropriate solution to the problem of operating multiple system images on a single hardware configuration has been announced. It is the Processor Resources/System Manager (PR/SM, pronounced "prism") announced on February 16, 1988. PR/SM provides hardware-controlled logical partitioning without any need for VM. With PR/SM, any 3090 E-model equipped with the feature will be able to run up to four independent operating systems per MP side. PR/SM sounds similar to Amdahl's Multiple Domain Facility (MDF) but with some improvements in dispatching capability and greater operational flexibility.

The remaining role for VM in MVS installations will be limited to running low volume test systems and supporting MVS production virtual machines on hardware not equipped for logical partitioning.

Lesson of the Legend

Let VM be VM. Keep it away from MVS.

See MVS Legends *page 104*

AUTOMATION
A Managerial Discussion

By Greg Coticchia

Automation. The word itself conjurs up many images from Kurt Vonnegut's 20-year-old *Player Piano* that portrayed a future society in which an elite few run the machines while the unemployable majority subsists on handouts in resentful idleness to the all-too-real world of the modern manufacturing plants where automation in the omnipresent form of robots crane, twist and weld American automobiles on an assembly line with no workers in sight.

The word automation has been with us since the early Greeks. *Automatous* is derived from the Greek words for self and move. Since that time, automation has often been viewed at the extremes. To automation and the technology that accompanies it has been given credit for most of the significant improvements in the condition of human life, especially in the workplace, and also most of the blame for the deterioration in related conditions.

Why? The existence of extreme views is not surprising when you consider how automation has permeated many aspects of daily life. Even what is considered automation changes as the technology that is so closely related to its application evolves ever so rapidly. Therefore viewing automation as simultaneously a benefactor and malefactor is not inconsistent. As automation has moved society from agrarian to industrial to what is commonly referred to as the present day post industrial phase, the bottom line of automation has not changed. It is still higher productivity. The driving force of business today is global competition. Technology in the form of automation, and therefore higher productivity, affects international competitive advantage since it has a role in determining both relative cost position and/or differentiation of a company's product or service — the two driving factors of competitive advantage.

Now automation is being applied to the automators. Those who originally brought us the information society and therefore the large data processing community of today are realizing that many of the "worker as machine operator" philosophies of the Industrial Revolution are still present in DP, despite the existence of technology's progress to the contrary. As competition grows internationally, it is not surprising that *automated operations* is the buzz word of today's VPs of MIS and data center managers.

Automated Operations

Lights Out? Unattended? Automated? For as many definitions and degrees of automation that exist there are probably twice as many for every definition in a particular industry. Therefore instead of adding yet another to confuse us, I suggest following the definitions put forward by Dr. Leilani Allen of Amdahl Corporation for three levels of automation in data center operations.

- Automated operations is the application of software to perform many of the routine operational tasks in computer centers with greater speed and accuracy than would be performed by humans alone. The goal here is to increase availability.
- Unattended operations is the application of intelligent hardware and software to provide a self-managing data center environment with minimal human intervention. The goal here is to reduce personnel costs.
- Lights out operations is the physical separation of data center staff from computer hardware. The goal here is to reduce facilities' costs.

As can be seen by the definitions above, automation is not just software, not just hardware, not just staff, but *all* of them! Implementing automation is not a "quick-fix" to any of these goals, no matter which one is correct for your data center. It will take a commitment and effort by your organization, reviewing not only product solutions but also how they work and where automation makes sense to improve productivity and quality.

This brings about another shade of the automation definition: every data center's goal in automation is not the same. As data center operations have moved from

Automated, unattended, lights out operations describe three levels of automation in data center operations.

The Power Of Automated Operations At A Single Point Of Control

Data centers serious about their plans for automated operations are beginning with the system console. Advance your own automated operations plan with AutoMate/MVS, a full function automation tool that solves many operational problems in today's complex MVS data centers.

AutoMate/MVS concentrates its power at the console to relieve operators from mundane, repetitive tasks and decrease the possibility of human error. The clear benefits are **increased productivity, improved system availability, and cost reductions** resulting from smoother, more reliable operations.

Improved Operator Environment

New, menu-driven operator interfaces offer RULES creation capabilities and easy-to-use ISPF display panels that provide organized, clutter-free screens for the operator.

All this can be done from a single terminal when AutoMate/MVS is integrated with other Duquesne Systems products. For a better operator workstation environment, use AutoMate/MVS with Multi-image Console (formerly SCON and GCD) for message consolidation; Remote Console/MVS (formerly ECS) for remote operator console capabilities; and TPX for session management.

NetView Compatible

As a tool for automating MVS console operations, AutoMate/MVS fills in where NetView* falls short by providing facilities for powerful message and command management. AutoMate/MVS utilizes an easy-to-use, rule-based system which eliminates the need to write and maintain complex assembler language MPF exit routines.

AutoMate/MVS can be an important step in your comprehensive plan for full system automation. For more information, call your Duquesne Systems account representative at **(800) 323-2600**, or **(412) 323-2600** in Pennsylvania.

**NetView is a product of* IBM, *Armonk*, NY

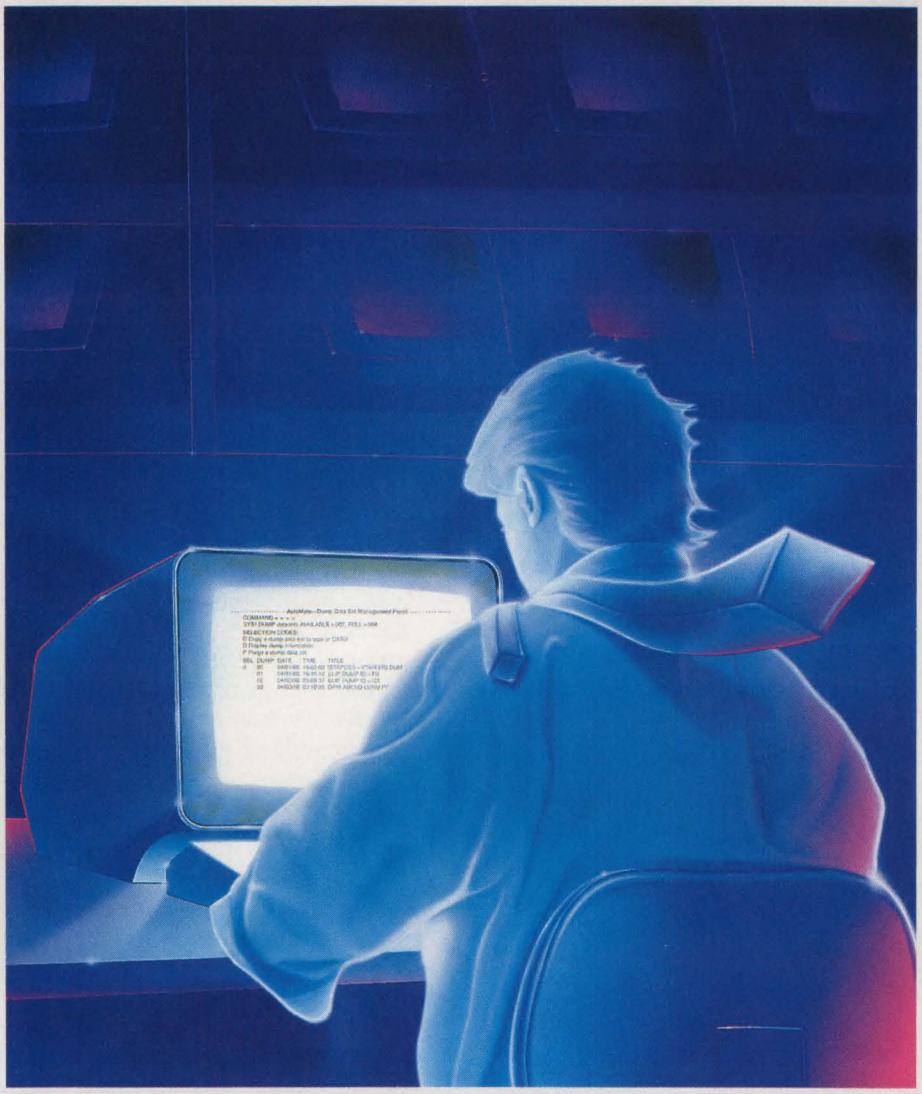

AutoMate/MVS integrates with other Duquesne Systems products to bring information quickly and easily to one single point of control.

AutoMate/MVS™
BY DUQUESNE SYSTEMS

DUQUESNE SYSTEMS
Two Allegheny Center
Pittsburgh, PA 15212
Telephone 412-323-2600

CIRCLE #1 on Reader Service Card ▲

truly manual operations to the current operator specialization (MCS, MTO, printer, tape drive and so on) of most shops today, great productivity increases were obtained. The same will be achieved through automation. However as expert system capabilities are imbedded into hardware and software for both guidance and real-time decision making capabilities, further productivity gains will be made. As seen in Figure 1, when mapped as productivity versus time in the far out reaches of time, unattended operations loom. What does this say to those of us who are committed and serious about implementing automation?

- Unattended operations may be or may not be a goal of your automation plans.
- Depite this, automation is the first step toward increased productivity even though not all the pieces of the technology puzzle are here to reach unattended operations.
- Most importantly automation is a process — not a product! There are few, if any, ''off the shelf'' solutions.

So how do you in your data center begin to justify and implement the automation process so as not to be left behind? And where do you begin? Choose the area with the highest payback and where tools exist today to get you started and take you into the next three-to-five years: the operator's console.

FIGURE 1

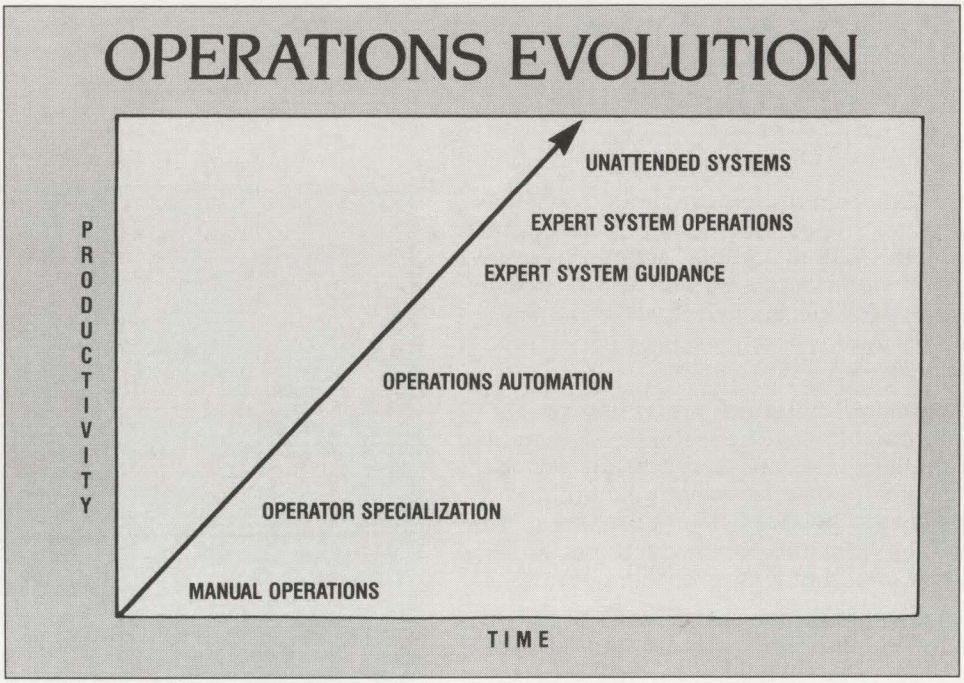

Console Operations Problems

Console operators face in the data center what many factory floor workers faced in their jobs in the 50s and 60s prior to the automation push in manufacturing — many dull, routine, boring tasks and little or no opportunity for career growth. Review the problems at the operator's console that affect availability.

Message Overload

The console operator is confronted with an overwhelming number of messages. In the average shop, operators are confronted with four to 10 messages a second! But simple message supression is not the answer. Controlling this traffic requires conditional logic for highlighting, low lighting, replying and rerouting messages to appropriate personnel and so on. It also requires, in some instances, additonal messages reworded or added for more meaningful control by operations staff (this is, of course, after message control is obtained).

Operators play an important role in this, the first part of laying the foundation for automation, since they are closest to what needs to be automated. Therefore whatever tool is implemented for this first phase of message management and simple automation should be operator oriented, requiring little or no programming skills.

Operator Error

No one likes to hear about mistakes on the job, but in many ways the operator's environment with its unfriendly user interface and ever-growing complexity of software and hardware becomes more difficult to manage. Unfortunately, system messages are presented to operators one line at a time. The operator spends time answering routine messages and cannot concentrate on more complex problems. But one axiom holds true: the faster the system, the faster it waits.

Many of these types of problems can be handled through simplification and presentation of operator information and commands. One solution is to use standard IBM software tools, such as Interactive System Productivity (ISPF), in conjunction with automation for simplified command entrance.

As noted here, all automation is not simply ''behind the scenes.'' Much of it is oriented to allowing the operator to become proactive at complex problem solving — not simply being reactive to system needs.

Many Awkward or Unknown Procedures

Again, operators are dependent on ineffective, incomplete, manual tools such as procedure manuals or product documentation. This is not to fault documentation on the whole. However when faced with the situations described above, there is little time to search correct procedures or documentation, even if it is available or current, for system solutions. Many operators are forced to reach for the phone for technical assistance and that robs the time and productivity of two people instead of one.

Difficult to Learn

Operations staff turnover due to present working conditions is not only inevitable but also a fact of life. Therefore getting appropriate payback from new personnel in the operations area can take years instead of months without new, friendlier tools added in addition to the application of message management with simple and complex automation.

This is a sampling of just some of the operations difficulties that affect current data center operations today. When can you begin addressing these problems? Now! In most cases you have probably already begun. Think about some products that have already started you along the automation process:

The Power of Automated Operations Is Here.

It takes power to automate. Not just iron and a little software. But a force behind the automation tool that can help when things get tough. Candle, the company that defined MVS performance with OMEGAMON,® has entered the automated operations field with AF/OPERATOR.™

Energize Just What You Need
Candle's family of automation products is flexible enough to meet your individual needs–from simple message management to remote operations. It's up and running in an hour, handling basic message traffic and doing flawless IPLs at warp speed.

Or Blast Off Into Full Automation
When you're ready to automate recovery or complex procedures, AF/OPERATOR enhances your growth. With an automated operations product from Candle, you're uniquely positioned to handle whatever you need in the engine room, including automated availability and performance when the time comes.

Streamline the Engine Room
Basic message management is the first step. And with AF/OPERATOR's sample library, that first step is a snap. Yet, your horizons are never limited because you have the power to tap into fully integrated automated operations with the ability to control MVS and its subsystems–IMS, DB2, CICS, and JES.

With High Performance
System overhead isn't an issue when you can have less than 1% CPU consumption. Nor is scheduling when you can make changes to the system on the fly, rather than waiting for an overnight 'gen.'

And Support From the Mother Ship
When you take a step into the dimension of automated operations, you want the best ground crew behind you. And Candle's support has received DataPro's highest rating. This support extends beyond round-the-clock customer service and education. Candle spends $30 million a year in R&D so you won't ever be left behind.

So whether you're going to the outer limits of automation or just streamlining the engine room, Candle is ready to join your crew. Let automated operations beam down into your data center. Call Terry Forbes today at **(800) 843-3970.**

Candle®

Copyright ©1988 Candle Corporation. All rights reserved.

CIRCLE #120 on Reader Service Card ▲

- Shared device products for DASD, tape drives and console consolidation and better, singular operations control
- SYSLOG management products for archival and retrieval of valuable SYSLOG information
- On-line monitors for performance reporting
- Schedulers to relieve operators from tedious, routine job scheduling actions
- Report distribution systems to control printing actions and routing
- DASD and tape management products for effective resource management.

The list goes on to include many homegrown solutions as well. So do not be afraid to build off your current operations productivity products.

Deciding What to Automate

As we have discussed, automation can be a never-ending process. Set your goals and objectives, both short term and long term, so you can effectively benchmark your automation progress along the automation odyssey. Be sure to form strategic alliances along the way because of the internal political resistance that automation usually receives. For example, if technical support management is charged with the automation project in your organization, be sure that operations management and even staff is included on project goals and product evaluation. Consulting operations is always beneficial since they really "own" the automation process.

When beginning the actual implementation, there are several steps and goals that should be included on your "to do" list.

■ *Review current operations actions*

This information is usually found on Syslog. A review of this information from archived datasets usually can highlight potential areas for automated console software usage.

■ *Begin with simple tasks*

Areas like message management or command simplification are options here. These should be able to be implemented with a simple rule-based language provided with optional menu-driven orientation for operations personnel.

■ *Move to simple automation*

Time driven auto-actions in concert with message reply automation for routine procedures also can be rule based. This eliminates many of the cases of operator error and action delays.

Simple automation should also allow for "mimicking" or intelligently imitating present operator actions. Your current staff probably makes great use of pad and pencil or stick-it notes to recall or remember past important alerts and messages. The automation tool you implement should also contain this capability to retain critical status information about jobs, dates, message text and so on.

■ *Expand to complex automation*

Complex automation languages are commonly based on strategic platforms such as IBM's TSO CLISTS and REXX (TSO/E V2). They will probably require programming skills but these should only account for 20 percent of the automation completed. But more importantly, common autotasks in TSO CLISTS or REXX should be able to be referenced from the

CICS AND VSAM SOFTWARE
$495 - $1,295
FREE TRIAL - OVER 6,000 SOLD

- ***OPEN/CLOSE CICS Files from Batch JCL - CICS/CEMT***
  Issue any CEMT SET command. ALLOCATE and UNALLOCATE files. ENABLE or DISABLE transactions. Supports up to 99 CICS systems on multiple CPUs. Purchase for $895 (OS, MVS) or $695 (DOS).
- ***SEND MESSAGES TO CICS TERMINALS - CICS/MESSAGE***
  Send messages to CICS terminals without affecting user's transactions. Saves the screen buffer, Tran-id, and Commarea and Displays the message. The user's screen and transaction are fully restored. Purchase for $695 (MVS or DOS).
- ***SEND COMPANY or DEPARTMENTAL NEWS - CICS MORNING NEWS***
  Automatically displays up to one full screen of news when the user signs on to CICS (CSSN), or logs on to CICS (CSGM). On line Help screens available. Full 19 X 79 screen available for each news item. Runs under CICS. No special hooks required. Purchase for $495 (MVS or DOS).
- ***CREATE CICS MAPS FAST - CICS/MAPR II***
  Includes a screen painting facility, full length data names for COBOL DSECTS, group-level occurs for repeated rows, and a COPY facility which allow users to build a new map using an existing map as a base. Purchase for $1,295 (MVS or DOS).
- ***FORWARD RECOVERY SYSTEM for CICS FILES - CICS/FRS***
  Supplies recovery programming for corrupted or lost VSAM files (KSDS, ESDS, RRDS). Flexible system which can select records based on variety of parameters and selectively update your backup files to get your users back on-line quickly. Purchase for $1,295 (MVS or DOS).
- ***ELIMINATE IDCAMS LISTCATS - LISTCAT PLUS***
  Replace bulky IDCAMS Listcats with LISTCAT PLUS. Lists VSAM entries in a one-line-per-dataset format and flags datasets needing attention. Includes all ICF files too. Makes tuning suggestions. Also prints a volume summary showing space utilization. Purchase for $695 (MVS) or $495 (DOS).
- ***EDIT, COPY, and DEFINE VSAM FILES from TSO/ISPF - IVU***
  Provides on-line access (through ISPF) to the most often used VSAM functions that would normally be done with IDCAMS batch processing. You can EDIT, DELETE, COPY RENAME, BROWSE, and INQUIRE on VSAM datasets without programming or writing JCL. Edit NON-VSAM files with record lengths greater than 255 bytes. Purchase ISPF/VSAM UTILITY for $1,295 (MVS or DOS).

MACKINNEY SYSTEMS
2674-A South Glenstone, Suite 112
(417) 882-8012

rule base for operator usage. Also, the TSO or REXX scripts will need to be supplemented with enhanced automation command processors, either written by your staff or provided by the automation package for more complete automation capabilities.

■ *Provide operator workstation capabilities*

It is clear that operators need better tools to do their job. Many facilities for automation, such as SYSLOG Recall, should come from the present day master console. As much as the console can improve, a better interface that builds on those facilities exists today.

The idea of a workstation environment for operators was only a pipedream a few years ago. By combining many useful software tools, the dream can come close to reality.

Providing a session manager and multi-image console product (for access to multiple MVS message streams) is a giant first step to consolidating and streamlining unnecessary message traffic from different consoles. This simplifies the operator's environment and makes automation gains more rapid. Beyond that, the session switching capabilities should allow access to SYSLOG analysis and review software, as well as ISPF menus for specific operator automated actions. These tools together provide the basis for a truly friendly operator workstation environment.

The future of this terminal is certainly spelled out in two words: personal computer. Automation solutions today deliver both internal mainframe and external personal computer solutions for early IPL support, external dial out capability, color and sound (audible alerts, voice support and so on). There is little doubt, as PCs proliferate, that this will serve the basis for future operator workstation capabilities.

■ *Take advantage of Open Software Architecture*

An effective operator automation tool for the future should be "open." It should be multi-vendor in terms of its interfacing. At one level, support should be implemented via current console messages and exits. At another, LU2 support is important in the short term for successful automation of panel applications such as performance monitors. This allows support across vendor products and the potential for unlimited application and growth of your automation.

Beyond this, remote operations capabilities are also important. This can be in two forms: remote display of the operator workstation/console and application-to-application interfacing through IBM protocols such as LU6.2.

Other Considerations

NetView and other IBM Standards

IBM is always to be considered when venturing into an area of commitment such as automation. Much of the future progress in this area will be delivered by IBM products — both hardware and software.

One of the most significant announcements affecting data center automation implementation was made in June, 1987. It was NetView. This is, of course, IBM's network management product for controlling and monitoring network performance. It also provides a procedural language, Network Communications and Control Facility (NCCF), that provides network operations a command list language similar to IBM's standard TSO facilities. Combined with several other NetView components, NetView provides a basic framework for future network management offerings by IBM. From the system operations side, it provides an additional tool for system/network automation.

For today's pressing need to address system operations, however, NetView relies on Message Processing Facility (MPF). MPF does three things well: message suppression, message suppression and message suppression. It is not a complete or easy-to-use facility. It also does not provide for operation staff involvement because the usual data center security system will restrict operators from accessing MPF. Therefore those closest to knowing and implementing the automation will not be involved. With 90 percent of the automation with NetView NCCF CLISTS a higher level interpretive language, the success of the operations automation project will be costly and slow moving from the start. This is where many Independent Software Vendors (ISVs) have aligned their products to take advantage of the time, effort and ease-of-use facilities, as well as lower overhead with a system console orientation versus a network management spinoff.

More importantly, the MPF/NCCF offering allows for no improvement that addresses the question of, "What have you done for your operator?" No menu-driven, combined-status displays for easy, one key stroke action commands are provided in the current Release 2 offering or are they planned.

IBM is also providing tools to support the system operations as well. With its announcement of REXX under TSO, many ISVs who have built their products on IBM standards such as TSO will readily take advantage of this and other IBM futures. This also applies to other IBM standards such as ISPF that was discussed earlier.

System vs. Product

When deciding on any product in the automation arena, differentiate between purchasing a "product" or a "system." The difference is not always obvious. Many systems are complex and therefore require higher levels of support and system integrators to install and tailor multi-component "products." This can leave you and your staff out of the automation loop and left behind with what was done two to three years down the road.

Products, on the other hand, provide a base starting point that allows you to grow and control the pace and success of your automation project. Again, multi-vendor solutions are normal implementing many products, both hardware and software.

Conclusion

There are really some simple realities in data center operations today to remember when starting your automated operations plans.

Automation is a process that you should begin *now* with multi-vendor solutions that are based on or are compatible with IBM's strategic directions.

This is how your data center can continue to grow and stay competitive in the global economy of the future that in many ways is here today. Gaining and keeping that advantage is not a short-term quick fix and neither is the underlying automation that will keep that edge alive. It requires, as other things do, a strong commitment. It also requires a seriousness to implement what is right for your data center and your organization's operations now and in the future.

ABOUT THE AUTHOR

Greg Coticchia is manager of the Operations Productivity Group for Duquesne Systems Inc., Two Allegheny Center, Pittsburgh, PA 15212, (412) 323-2600.

Is your PRODUCTION SCHEDULE out of control?

Do RERUNS have YOU on the run?

Perhaps you should consider . . .

CONTROL-M

The Automated Production Control System that puts YOU in control.

- Totally automates Job and Event Scheduling
- Improves system throughput
- Significantly reduces reruns:
  - Ensures jobs run in the proper sequence
  - Submits Jobs ONLY if ALL prerequisites (manual, predecessor jobs, datasets, hardware resources etc.) are met.
  - Automatically updates JCL and Control Cards, eliminating manual errors.
- Provides on-line management visibility and control
- Schedules both batch jobs and Started Tasks
- Automatically opens and closes CICS files
- Automates restart and recovery
- Simulates hardware and software changes
- Puts job documentation on-line
- Produces immediate notification of problems
- Installs in 1 hour WITHOUT ANY SYSTEM HOOKS
- Makes schedule definition quick and easy

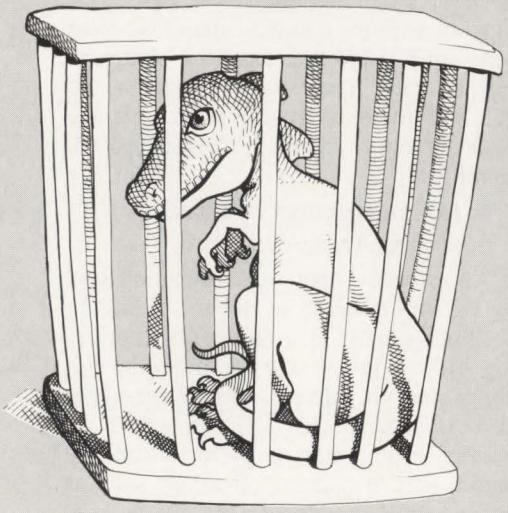

CONTROL-M tames the RERUN monster

MVS, MVS/XA, JES2, JES3

TONE Software Corp.

1735 South Brookhurst, Anaheim, CA 92804-6491
(714) 991-9460 • Telex 4974583 • 1-800-833-8663

CIRCLE #144 on Reader Service Card ▲

Strategic Alliance

Landmark Systems and Morino Associates Demonstrate Advantages

By Fred Newberry

At first glance, Landmark Systems Corporation and Morino Associates, Inc. have little in common. True, both are located in Vienna, VA and both are driven in an entrepreneurial style by technician-managers who still code a few lines now and then. Customers rate both vendors exceptionally high in terms of technical support. Both of their flagship products, Landmark's The Monitor for CICS and MICS from Morino, are stout, well-behaved products that may be installed without regenerating the entire Western Hemisphere.

However there *are* differences. Morino Associates is an older company, larger by every measure. Landmark is privately held; Morino is a public company. Landmark is still principally a single-product company although its introduction of Eyewitness earlier this year will change that characterization. Morino offers dozens of integrated products in the MICS family. The personal and managerial styles of their CEOs occasionally contrast dramatically.

To the advantage of their customers, however, the similarities clearly outweigh the differences. In 1987, the companies announced a unique *strategic alliance* agreement that saw the development of an interface between Landmark's The Monitor for CICS and Morino's MICS CICS.

The interface allows MICS CICS users to directly input data from The Monitor. Without the interface, users who wanted to feed data from The Monitor to MICS CICS would first have to convert it to PAII or CMF formats. The interface, then, provides three major benefits to users. First, it bypasses the need to convert formats. Second, by eliminating the need for CMF itself, CICS overhead is significantly reduced. Third, by processing The Monitor data directly, all data fields can be summarized in MICS, promoting more comprehensive analysis and capacity planning. In summary, the interface provides users of MICS and The Monitor with a superior combination of tools for troubleshooting and capacity planning of CICS systems.

When Landmark Systems began selling The Monitor for CICS in 1983, first-year revenues were a mere $12,000. By the end of 1987, Landmark had boosted its revenues to $13 million and it now boasts more than 3,000 installations of The Monitor. Without outside financing the company's founders, Patrick McGettigan (currently CEO) and Katherine Clarke (currently VP of Technical Support), both quit their jobs to start the company.

The Monitor for CICS has its origins at Blue Cross/Blue Shield in Washington, DC where a monitoring system was developed in the late 1970s to gain more control over response time. The resulting system met with a great deal of enthusiasm, leading the developers to ready it for market. Landmark Systems released The Monitor in mid-1983 and by mid-1987 more than 2,500 DP sites were licensed to use the product worldwide.

Founded in 1973, Morino Associates markets an integrated line of software products to control costs and improve the performance of MVS computer systems. Its flagship product is MICS I/S Management Support System, a family of modules that generates information useful for management alert reporting and cost management, network management, capacity management and more. Some 5,100 copies of its products are installed at more than 1,200 sites in 25 countries. Net revenues for 1987 were $33.3 million.

The MICS CICS Analysis module provides CICS reporting capabilities for system availability, service levels, transaction throughput, resource usage and transaction profile analysis. It provides supporting information for cost recovery of DB/DC resources and for capacity planning. It maintains measurement data from IBM's PAII and CMF products and Landmark's The Monitor for CICS.

CICS monitoring is critical if such areas as on-line response time, problem determination and resolution, application debugging and tuning and data management are to be controlled. A CICS monitor provides the system manager with answers to such real-time problems as: What is wrong? How can I fix it? How can I anticipate the problem before it happens again? In detail, its capabilities are as complex as CICS itself.

For system managers who want to minimize CICS problems, a real-time task activity monitor is needed to pinpoint problem tasks as they occur. Aberrant tasks can then be quickly canceled before they degrade system performance or worse, cause CICS to crash. For application tasks, a CICS monitor should show task status and identify such problems as enqueue lockouts, VSAM string waits and high resource consumption.

The Monitor/MICS interfaced is currently installed in approximately 20 installations. At one of the installations, Joe Bell, project manager for Hallmark Cards (Kansas City, MO) reports that great improvements in capacity planning were made possible by the interface. A user of MICS and IBM's CMF since 1982, Hallmark Cards wanted to convert to The Monitor for two reasons. First, it desired an on-line monitoring capability. Second, it suspected that it could save on overhead. According to Bell, Hallmark subsequently verified that overhead savings were possible.

Without the interface, Hallmark would have had to concurrently operate both The Monitor and CMF if it wanted on-line monitoring capabilities. "We would have had to justify The Monitor by its real-time capabilities alone and absorb the additional overhead required by operating two monitors. I do not think that would have been too easy," Bell says.

A 17-year veteran of the software industry, Bell is gratified by this display of cooperation. "I'd like to see more of it. The users would obviously benefit and the software vendors would benefit, too."

At San Francisco-based Levi Strauss & Company, Technical Consultant Phil Horvath has used MICS and The Monitor since late 1987. Performance and service level agreements and the chargeback system rely on MICS, which in turn uses input data from The Monitor. The interface makes it all possible.

Horvath explains, "It takes data from The Monitor and produces records that MICS is able to read and manipulate for performance analysis and chargeback. Without it, we would have to go to CMF or some other CICS monitor to extract the data. The Monitor seems to provide a little better performance with less overhead. So without it, the monitor we might have to select would be less than optimum for our operation."

Horvath welcomes the increasing levels of cooperation between IBM and independent vendors and the strategic alliance between independents themselves. "Nowadays, vendors have to be able to support, to some extent, other vendor's products," he says, pointing to the interface that VM Software, Inc. (Reston, VA) created to the UCCEL UCC-1 tape management system now controlled by Computer Associates.

At Landmark Systems, Robert Lewis, development manager for The Monitor,

An Interview with Patrick McGettigan

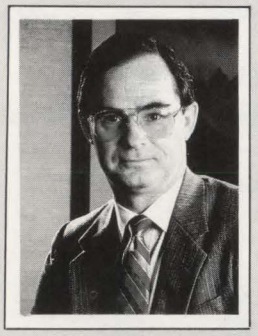

Patrick McGettigan, CEO of Landmark Systems Corporation, received Washington's "Entrepreneur of the Year in High Technology" award from Arthur Young and VENTURE magazine in June, 1988. He developed The Monitor for CICS while he was at Blue Cross/Blue Shield. McGettigan and his partner, Katherine Clarke, founded Landmark Systems in 1982; The Monitor was introduced the following year. He was the company's technical leader until 1984 when he moved into executive management.

MAINFRAME JOURNAL: *Is this agreement with Morino Associates the first for Landmark?*

McGettigan: Yes. Now, The Monitor supports six 4GLs and as many databases, but it has cost us too much because we do not have the simple relationships we need.

MJ: *Examples of strategic alliances are rare. What are the obstacles to such agreements?*

McGettigan: Some time ago I went to a session on strategic alliances. It was an abominable presentation. What the panelists really demonstrated was that strategic alliances are done the hard way, from an adversarial point of view. There is so much fear and paranoia in this area, that most vendors do not practice a deliberate attempt to make strategic alliances.

MJ: *"Paranoia?" Can you explain?*

McGettigan: Well, I can understand some of these fears. Partners in a strategic alliance today can easily become competitors tomorrow. Vendors have fears about how a commitment might tie their hands or impede them down the road if a competitive situation should develop.

MJ: *How about benefits? What is the main incentive to such agreements?*

McGettigan: Customers always provide the leverage that creates momentum for strategic agreements. There are a couple of directions you can come from. First, the technical side. It is important to have a direct line of communications with firms like IBM, Cincom or ADR to provide us with the limited information we need. The second side to a strategic alliance is business and marketing. There are ways this can be done that do not cost you anything such as seminars, cooperative advertising, that sort of thing.

MJ: *Briefly describe your management style.*

McGettigan: I'm a bleeding heart people-oriented person. My partners and I work hard at enrichment of the work environment and our employees' lives in general. I believe it has a lot to do with how stable we are. In general, I believe if it feels good, do it.

has a unique perspective on the development of The Monitor/MICS interface. Until he joined Landmark in 1987, he had development responsibilities at Morino. He recalls the initial meetings in which the specifications for the interface were first discussed and specific technical challenges first addressed.

"We encountered a number of problems in the initial design of the interface," he recalls. One problem was that The Monitor data could come in two formats: compressed or uncompressed. One challenge was how to provide both formats to SAS, the reporting language favored by MICS. Eventually the problem was solved by writing an Assembler front end to process the data before passing it on to SAS.

Another problem was solved by modifying code for The Monitor so that it produced interval systems records. The problem was that MICS is interval-driven while The Monitor is more activity-driven. In other words, MICS is keen on summarizing by the hour. If MICS does not see any data for the hour or the specified interval, it might lose other data generated by The Monitor. To prevent this Landmark provided a modification that would force system information out every hour even if there were no activity in that interval.

Carl Meredith, currently development advisor, has been with Morino Associates for eight years. He sees The Monitor/MICS interface as increasing the return on investment of MICS. "If a shop relies on MICS, it becomes a critical part of its operations. Our users do not want to be held back by not being able to input critical data to MICS."

Meredith agrees that the kind of cooperation that has led to The Monitor/MICS interface is not normally found in the software industry. "We had some advantages. Landmark is a very concentrated effort; there is basically no overlap in functionality that allows competition between our two companies." He is happy with the results and welcomes the benefits of cooperating with another vendor. "You can always get a better view standing on someone else's shoulders."

He sees benefits all around. "Morino Associates has always tried in the MICS product to deal with as many popular data sources as possible. The advantage to us is an increased prospect base. With the interface, we now have analytic components that will analyze performance of MVS and perform CICS capacity planning, accounting and chargeback using data from either CMF or The Monitor. Consider sites that use The Monitor. In such sites, we would not otherwise be a candidate for selling CICS component or CICS billing."

A similar argument is made for Landmark. The interface gives Monitor users another reporting tool to get at transaction data for capacity planning and modeling.

But to Meredith, it is the user base for both products that stands to benefit the most. He sees both visible and invisible benefits. "The visible benefit is that the interface allows users to install the monitoring tools they consider best for their shops. In this way, they tend to end up with fewer errors and less overhead." The invisible benefit, he explains, is that

XA-RELO / XR-RELO

CICS Performance Optimizer

As you may have already discovered, converting to XA does not necessarily mean your CICS performance and storage problems are over. Now it's time to let the powerful features of XA-RELO provide the solutions. (XR-RELO provides similar solutions for MVS/SP).

- **Enjoy faster response times**
- **Improve internal throughput**
- **Eliminate all CICS storage compressions**
- **Eliminate virtual storage constraints**
- **Eliminate Short-on-Storage conditions**
- **Increase the Dynamic Storage Area (DSA)**
- **Eliminate all program fetches from the CICS load library during execution**
- **Reduce system I/O and WAIT time**
- **Allow all programs and mapsets to reside in the XA address space without any recompiles or modifications, including macro level programs**
- **Easy to install ... less than 30 minutes without any system modifications or program changes**

— Call now for a free trial —

(800) 542-7760 • FAX (205) 833-8746

Quantum International Corporation

"Superior Solutions"

CIRCLE #63 on Reader Service Card ▲

An Interview with Mario Morino

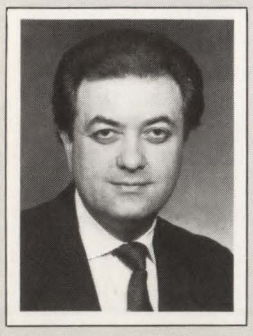

Mario M. Morino is CEO and chairman of the board of Morino Associates. He began his career in applications systems with General Motors and gained management experience with Eaton Manufacturing, U.S. Time Sharing and the Bureau of Naval Personnel. He subsequently became director of program products for PACE Applied Technology. The entrepreneurial bug hit him in 1972 and Morino Associates was formed a year later. The company introduced its principal product, MICS, in 1979.

MAINFRAME JOURNAL: *You have been an outspoken advocate of strategic alliance agreements for many years? Why is the software industry still lagging in this area?*

Morino: Such agreements are increasingly essential yet most vendors are brought into it dragging and kicking. It is hard because many still take a parochial view of their products. The fact of the matter is that users today exist in an intensified multi-vendor environment. And that means all vendors will have to work together more, including IBM.

MJ: *Is this the first such agreement for Morino Associates?*

Morino: Not at all. We have various levels of agreements. In addition to Landmark, we have formal situations with companies like BGS Systems and Amdahl. We are also working on a more informal basis with Boole and Babbage and Duquesne Systems. We are currently talking to Candle Corporation.

MJ: *What is the incentive for such agreements?*

Morino: It is very much due to customer demand. If we think a product is going to be a winner, we take on the support for an interface ourselves. MICS is smack dab in the middle of things. We just have to provide integration. Other vendors develop interfaces to MICS themselves and that is good for us.

MJ: *How were the costs and expenses of this development effort apportioned between the two companies?*

Morino: The two companies shared the knowledge. The key part of these agreements is not the development effort but technical currency. We typically take on the development costs ourselves. The key is to get the other company to share technical information up front.

MJ: *How do you plan to recover those costs?*

Morino: We assume we will recoup our costs because of the fact that Landmark's The Monitor for CICS is such a successful product that it will help us sell. When we support a product like The Monitor, we are convinced it is going to be a mainstay product with high penetration in the industry. That is good for us.

MJ: *Would you briefly describe your management style?*

Morino: I am fairly demanding. I am very much driven to a heavy client-orientation. I am an old fashioned in-the-field type person. I like to be out there with the users. That is where the action is. That said, I am very informal. We have a typical high-tech environment: high expectations of high-quality people working in a loose organization.

because Landmark and Morino Associates are constantly talking to each other, users will not get surprised by incompatibilities.

Landmark's Lewis takes it one step further. For him, the biggest advantage of the relationship is that it broadens the information base of both companies. "We generally talk only to CICS systems programmers; Morino generally talks only to performance analysts and capacity planning people. Via this interface, we share the information and get feedback from people we normally would not talk to." That is a rich benefit, indeed.

ABOUT THE AUTHOR

Fred Newberry is a free-lance writer specializing in high technology and computer-related topics.

Fundamentally, SQL offers users freedom of choice. To the extent all RDBMSs migrate toward SQL as their common access language, it will encourage compatibility at the application level. What is under the covers of the database structure will be transparent to the user. Applications developed under SQL/DS, DB2, CYGNET, SUPRA (Cincom Systems, Cincinnati, OH) or others can extract information from common databases. This capability holds out great promise for successful operations in a decentralized mode.

Finally, SQL promises compatibility with IBM's future direction, a factor no company can ignore. Since IBM has identified SQL as the key database sub-language in its Systems Application Architecture (SAA), customers who use a DBMS that supports SQL can rest assured they are not deviating into an area that would prevent them from taking advantage of future IBM developments.

Conclusion

Getting the most out of SQL/DS and the relational model requires that DOS/VSE shops adopt a new application development mindset. To the extent that they have ad hoc methods for requirements analysis and project management, these shops should adopt formalized processes for both developing standards and managing the application development life cycle. They should also emphasize communications between users and MIS staff and internal communications among the DP staffers.

The very flexibility of SQL/DS gives unaware users just enough rope to hang themselves. When data is not managed properly, a very effective tool is potentially more damaging than a less effective tool. Used improperly, SQL/DS may become a maintenance nightmare, as users proliferate applications. A DBMS by itself is not capable of identifying data resource usage. Only a data dictionary can provide this information.

To work successfully with SQL/DS, data centers must be clear about some SQL/DS misconceptions (see Side Bar titled, "Five Misconceptions about SQL/DS . . . Or Are They?"). They must also adopt some specific assumptions about applications development and RDBMS performance. These assumptions include:

- Acceptable performance demands a staff thoroughly trained in the use of SQL/DS and dedicated to sound system design and constant tuning

SQL/DS *from page 71*

- Using SQL/DS to build isolated tables gives you some incremental benefits but you pay enormous performance penalties
- High performance is frequently a trade-off with ease of use and installation
- A shared database environment for high performance update applications and ad hoc reporting applications will not be satisfactory for some users
- Without a management process in place to support integrated performance-based design, SQL/DS will result in very limited functionality.

SQL/DS and the relational model emerged because aggressive users were driving traditional software beyond its design limits. While the traditional approaches sufficed to automate the back office, the majority of front-office applications still are relatively primitive in terms of automation. This is the promise of RDBMS technology: it will move the sophisticated, on-line computer environments that for 20 years have supported production work to the desk of the customer service representative, the sales person on the road and the corporate boardroom.

ABOUT THE AUTHORS

Robert D. Smith, founder and president of Sysmith, has more than 15 years experience in a wide spectrum of software development areas. In addition to serving as director of technical support for Culinet Software, his background includes experience as a management consultant, author and speaker. Smith holds patents in the area of value-based security. Founded in 1984, Sysmith provides a range of professional services in capacity and facility planning. Sysmith, Inc. 13763 Fiji Way, Bldg. E, Marina del Rey, CA 90292-6909, (213) 822-7557.

Dr. Lani Spund is a senior management consultant on the Sysmith staff. He is a pioneer in the development of structured design, programming techniques and database management systems.

For information on advertised products, fill out the READER SERVICE CARD elsewhere in this magazine.

DASD Growth *from page 44*

activities outlined in the package. The most common of these activities is allocation of new production files.

The final step in processing a data management package is for the data manager to review the package and verify its status as a budgeted project. If the package passes the tests, the data manager allows the process to continue. The next step in the process produces the JCL necessary to accomplish the required file management. ADMS then places the package on an archive file and the resulting use of DASD is logged into the report generation file. From this file, the quarterly Forecast vs. Actual report is generated. Further, the performance/capacity unit is developing a MICS-based facility to locate candidate volumes for allocation or movement of files to be placed on any pool. The facility will use space and performance requirements to develop the volume candidate list. This facility is necessary for the planned automation of the data management component of the automated data management system. When this process is perfected, the role of the data manager will then be to review request packages and allow the files to be allocated/reallocated/deleted or reject the package and send it back to the requestor for modification.

Summary

Storage management and the resulting architecture for controlling use of storage resources is intended primarily to control the use of a rapidly expanding resource and ensure efficient use of the resource. Imposition of a storage management structure need not cause lowering of productivity in areas required to follow the rules necessary for the success of the effort. Connecticut National Bank has developed a storage/data management structure that assures both adherence to the rules of effective storage management, rapid turnaround of DASD and data management requests and "just in time" purchasing of data storage hardware.

ABOUT THE AUTHOR

Steve Adil does storage management strategic planning for Connecticut National Bank, MSN 348, 150 Windsor St., Hartford, CT 06120, (203) 722-9330.

XA-VSAM / XR-VSAM

Improve VSAM performance as much as 90%

- **Cuts hours off daily production time**
- **Significantly reduces I/O and WAIT time**
- **Analyzes and dynamically tunes the performance characteristics of all VSAM datasets**
- **Automatically provides optimum VSAM buffer management for maximum efficiency**
- **Eliminates dedicated resources now used to tune VSAM datasets**
- **Requires no program or system modifications**
- **Easy to install . . . less than 30 minutes**

— Call now for a free trial —

(800) 542-7760 • FAX (205) 833-8746

Quantum International Corporation

"Superior Solutions"

VMOPERATOR

Putting Muscle Into the VM Operator Console

By John Kador

The weakest point in many VM Data Centers is the VM operator console. As VM networks become faster and more elaborate, the probability rises that important messages will not get timely responses, or worse, be ignored altogether. At the same time, the operator console in the data center becomes a bottleneck as it is the sole place from which commands may be issued or problems diagnosed.

VMOPERATOR from VM Software, Inc. (Reston, VA) is a management system for the operator console. It enhances the interface between VM and the DP operations staff, provides full-screen and extended color control capabilities, automatic message handling and simplified on-line review of the operator console. Critical messages that previously slipped by unnoticed can now be held on the screen and displayed in an eye-catching color. Also, VM problems can be easily traced by systems programmers from their own terminals.

Here is a paradox. The faster a VM machine is, the faster it waits. Consider the increasing rate at which messages arrive. Most messages are informational, requiring no specific action. But they distract the operators from the remaining messages, such as tape mount requests that need timely operator interaction. If these messages are highlighted and appropriately directed, 75 percent of the wait time can be eliminated. Positive benefits, in addition to faster tape mounts, come from more rapid message response during initialization and shutdown. These are the benefits promised by VMOPERATOR.

Message Rate Relief

For important events such as tape mount requests, it is important that such messages be highlighted on the screen until the operator acts on them and manually deletes them. Some messages (such as INT REQ messages) are held on the screen and removed by VMOPERATOR when the condition is corrected. The routing tables supplied with the product includes handling for critical but easy-to-miss messages such as the one indicating that spool space is full, a condition that if not corrected can crash VM.

Console message suppression is one of the major applications of VMOPERATOR. If data center operators do not get some relief soon, they will become the inevitable victims of sensory overload as they are bombarded by increasing volumes of console messages. Message rate relief also provides an uncluttered console that, in turn, results in reduced wait times as operators are less distracted from acting on the remaining messages.

Without relief, the data centers they manage will increasingly resemble the air traffic control situation we have today with delays everywhere, occasional crashes, more near-misses, rising costs and burned-out workers.

More than 90 percent of the messages issued by operating systems are routine to the point that they should be handled without operator intervention. If the messages flashing over the operator console were segmented into a pyramid, the bottom 60 percent constitute informational messages that require no operation reaction. Examples include the User ID logon messages. Messages requiring pre-established trivial responses (''always or never'' situations) represent another 25 percent of the total. These messages require system intelligence to process and eliminate. The remaining 15 percent require a decision; something approximating human intelligence to handle properly. The responses to these items depend on variable environmental or operational conditions.

VMOPERATOR routing tables examine messages coming in to the VMOPERATOR system and contain instructions on actions to take for specified messages. The actions taken can include:

- Display the message on the operator console
- Ignore the message (although it is logged)
- Act on it (usually execute an EXEC routine).

On-line Console Review

No one likes to search through piles of printed logs to determine such things as when a file was printed. That is why VMOPERATOR records all message traffic in a log file for review. When users ask questions about past events, the answers are immediate. On-line, the reviewer simply scrolls back and forth through the current system log or searches the logs from previous days. Unlike spooled console logs, this log can be reviewed concurrently by several users. The on-line log also shows all current activity. The operator needs to take no action (such as closing a console spool) to allow other users to examine the log. On-line console review capabilities include being able to scroll forward and back through the log, shift the display left or right, search for a specified character string and extract all lines contained in a specified character string for further examination or printing.

Another feature of VMOPERATOR is that it allows authorized users to use their own terminals as an extension of the operator console. Analysts can then monitor the current system status, review past messages and invoke operator commands as if they were being entered directly into the operator console. This allows a systems programmer to connect to the operator console without going to the data center. If systems programmers have a terminal at home, they can obviously diagnose and fix problems without coming in.

Systems programmers often want to issue operator commands from their own terminals rather than the system console. Normally, however, the system console log shows only commands issued by the system operator. For command auditing, this is a problem. For example, if a systems programmer has forced a user off the system, the VM system console log shows only that the user was forced off but does not identify the systems programmer. With VMOPERATOR and VM/SP Release 4.0 or later, a site can audit the use of these commands, aiding in problem tracking.

Do you ever wonder how long it takes your data center to mount a tape? VMOPERATOR tracks the amount of time required for a tape mount and keeps this information in the system console log. Tape mount times can be easily extracted for the log and displayed.

Creative Programmer

Users of the VM operating system tend to be creative people. As it comes from IBM, VM is very limited. Luckily, VM is also very flexible. Unlike MVS, which

VMOPERATOR
VM Software, Inc.

does not allow users much opportunity to play around, VM offers creative systems programmers a great deal of latitude. Anyone who has worked with VM for an extended period of time has been forced to become creative to get the most out of the system.

Jennie Bohart, systems software specialist at Trane Dealer Products Group, is a model of this kind of creativity. Bohart, case in point, is primarily responsible for the VM system that consists of an IBM 4341 P12 running VM/SP plus two 9370s, a Model 20 and a Model 40. On the software side, Trane runs many of the products from VM Software including VMBACKUP, VMSCHEDULER, VMSECURE, VMSPOOL, VMBATCH and VMCENTER II. It also runs FOCUS (Information Builders, NY) and SQL on three database machines. The Tyler, TX company manufactures air conditioners.

VMOPERATOR was installed in late 1986 primarily on the strength of Trane's experience with other VM Software products. "We were impressed by how clean the code was and the sound interfaces between all the products," Bohart says.

At Trane, VMOPERATOR is used primarily to manage the operator's console. In the past, Trane's console operators were used to dealing with the MVS system and did not have a lot of opportunity to interface with VM. The primary VM interface is with tape mounts. MVS, of course, does not like to have a scrolling console screen, so the operators had a hard time remembering to clear the screens on the operator screen. As a result, whenever someone wanted a tape mount, they often had to call the computer room to tell the operators to find the tape mount messages that had the additional handicap of not being highlighted.

At Trane, VMOPERATOR highlights and holds tape mount messages and system problem messages. "VMOPERATOR gives me the ability to keep a lot of the trash messages from showing up on console. It writes them out to a log file so I can inspect them later," Bohart says.

What kinds of messages? "We commonly suppress User ID logon messages. In the normal course of attaching and detaching tape drives from User IDs, VMTAPE spits out a lot of extraneous messages that we roll off. Also, there are a lot of messages from NCCF/VTAM that our network operators handle that VM operators do not need to see. In summary, it results in less confusion for the operators," Bohart explains.

Trane also uses VMOPERATOR to manage the SQL database machines. It has several users who need the ability to issue operator-type commands. Trane gives these users access to VMYIAMOP, a function that provides tailorable screens, akin to an operator console. With proper authorization, this screen gives users the ability to issue most of the commands the operator can issue without having to call the console operator.

As an example of Bohart's creativity, she uses VMOPERATOR to issue VMTAPE mount requests from the database machines. She explains, "SQL is not friendly when asking for tape mounts when it comes time to do an archive. We trap those messages to issue a VMTAPE mount request. This has been a big help. I do not know how we could have done it otherwise." This function is not documented in the VMOPERATOR manual.

NEW

VM PERFORMANCE GAINS WITHOUT THE PAIN.

VM1 QuickCopy is the painless way to boost performance on your VM system. By dramatically speeding up all Copyfile operations – one of the most resource-intensive functions in VM – QuickCopy makes lightning copies and improves the performance of DIRMAINT, PROFS, Focus, SAS and other applications. Your entire system will work more efficiently, and so will your systems staff, programmers and end users.

QuickCopy fully supports all of VM's standard Copyfile commands and functions, so it "drops" right into your system. There's no training, no modifications to the system, and no headaches. QuickCopy even installs quickly – under ten minutes in most cases. And it's priced to fit painlessly into your budget.

See for yourself why so many other VM sites are achieving performance gains without any pain.

Call now for more information, or circle number 147 on the Reader Service Card.

1-800-321-4VM1
(in CA, 213-641-2600)

VM1

The Leader in High Speed Copying.

SAS is a trademark of SAS Institute • Focus is a trademark of Information Builders

DOS, OS, or CICS Frustration?

BIM gets it out of your system.

BIM presents a line of proven programs that maximize your system's capabilities, saving you time, labor and expense. These program products help get the most out of your system and people.

BIM-VIO — DOS/VSE Virtual Disk Drive. Moves the Standard Label Area directly into memory and allows for other heavily used non-permanent files to be moved into memory as well. NEW

BIM-PACK — Automatically compresses selected VSAM files transparent to applications and end users under DOS. NEW

BIMWNDOW — Multiple terminal sessions concurrently at CRT under **DOS** or **OS** VTAM.

BIM-EDIT/DOS — The most powerful, flexible full screen editor available for DOS/VSE.

BIM-EDIT/MVS — All of the features of our popular DOS editor and does not require the overhead of TSO. Can be accessed directly from VTAM or from CICS or other terminal subsystems. NEW

BIMSPOOL — Prints output in POWER/VSE spooling queue on local or remote 3270 terminal printers. (Received ICP Million Dollar Award 1982).

BIMSPLSR — Optional laser printer support for BIMSPOOL.

BIMSPOON — On-Line to Batch Print Spooling. Prints data passed from CICS application programs into the POWER spooling queue. NEW

BIMSPLIT — May be used separately or with BIMSPOOL to print parts of an existing job to terminal printers at separate sites. NEW

BIM-PDQ — POWER Dynamic Queuing performance enhancement. Eliminates 85% of the I/O to heavily used POWER queue.

BIM-PADS — Automatically alters or deletes DOS POWER spooled job entries at preset intervals. NEW

BIM-ODIS — Comprehensive problem analysis and display of operational CICS system. **ODISTRAK** is an optional historical reporting feature to be used with BIM-ODIS to generate reports relating to system usage. **DOS** and **OS.** NEW

BIM-BUFF — Significantly increases the performance of VSAM under DOS by dynamically managing VSAM buffers. NEW

BIMTEXT — Word processing, document composition system. Create formatted documents from free-form input. **DOS** and **OS.**

BIMSWAP — Switch local 3270 BTAM terminals between multiple CICS partitions without special hardware or additional ports.

BIMCMPRS — CICS 3270 data compression system. Reduces response time for remote terminals significantly. **DOS** and **OS.**

BIM-FMAP — CICS BMS on-line map generation and maintenance. **DOS** and **OS.**

BIMECHO — Copies one CRT's output to another or printer for problem determination and demonstration. **DOS** and **OS.**

BIMP3270 — Comprehensive CRT screen image print facility. Copy to terminal printers or spool queue for system printer. **DOS** and **OS.**

BIMSERV — On-line display of library directories and entries, VSAM Catalog entries, disk VTOC's, etc.

BIMCNSOL — Multiple/Remote System Console function for CICS. Display-only or full input/display versions available.

BIMMONTR — DOS/VSE System Status, Performance Measurement, and POWER Queue display.

BIMSUBMT — On-line Job Edit and Submission facility.

BIM programs are cost-efficient, some less than $800, average $2500. You can save even more with our group package offerings. Products are available on permanent, annual, or monthly licenses, and shipped on a 30-day free trial basis. Product documentation is available on request.

BIM also performs systems programming consulting, with consultants based in Minneapolis and Washington, D.C. Computer time services are also available on our 4331-2 system, on-site or remote.

BIM

B I MOYLE ASSOCIATES, INC.
5788 Lincoln Drive
Minneapolis, MN 55436

612-933-2885
Telex 297 893 (BIM UR)
Member Independent Computer Consultants Assn.

Issue Operator Commands

Trane uses VMOPERATOR to issue operator commands outside the IMOP screen. It uses this function to allow its disconnected service machines to do things like kick off an SQL archive. One of the machines is autologged by VMSCHEDULE and issues an operator command to tell the SQL database machine it is time to start the archive. In the past, Trane used to do SQL displays just to keep them for historical purposes in the console log.

She considers it somewhat of a limitation that VMOPERATOR does not permit her to restrict some commands when, for example, issuing VMSCHEDULE or VMTAPE commands. Doing so allows users to access all commands available to the operator's User ID, powerful commands such as CMS, CP, End or Abend usually not available to IMOP users. ''I'd hate to kick off my full-volume backups in the middle of the day inadvertently,'' Bohart laughs.

Part of the problem is that Trane shares tape drives between the VM system and two MVS systems. It uses the Shared Tape Allocation Manager from Duquesne Systems (Pittsburgh, PA) for which VM Software wrote an interface to VMTAPE. VMTAPE allows Trane to allocate tape drives to a user and bypass any kind of security. ''I'd as soon everyone not be able to do that,'' she adds.

''Using VMOPERATOR, we have built real friendliness into our VM system for handling printing facilities. We have written some nice EXECs to make it easier for the operators. I cannot imagine running VM without VMOPERATOR. It has made VM real usable for the operators,'' she concludes.

Missed Important Messages

Al Columbia is senior systems programmer at Pfizer, Inc. He works in Groton, CT in the New York-based pharmaceutical company's Corporate Information Services data center. The center operates an IBM 3083 under VM/HPO. It installed VMOPERATOR in mid-1986 on the strength of its experience with VM Software, Inc. The initial justification was to hold messages on the console screen. ''We constantly missed important messages to us. It got to the point where we had to tell users, 'If you want a tape mount, call us,' '' Columbia recalls. In addition to holding messages, he soon found how beneficial it was to filter out messages he did not want operators to see.

See VMOPERATOR *page 113*

DB2 *from page 66*

up the overall system logging rate (the logging will be eliminated in Version 2.1). It will always cause high CPU utilization due to high repetitive scan rates of the buffer pool(s) and inserts of intermediate results into temporary work tables.

- Returning a greater number of columns during a query will have a noticeable impact upon elapsed time and CPU consumption. Combining similar (repeating field) data into one large column will reduce both elapsed times and CPU requirements. Representative timings are illustrated in Figure 4.

Performance Considerations for Application Design

These considerations are presented as general information and may or may not be applicable to any specific project.

- Always use 'close = no' for tablespaces that will be used on-line. This must be specifically coded when the tablespace is defined because the default option is 'yes'. The elapsed time to open a tablespace has been measured to require up to five seconds. The close can require an equal amount of time. Obviously, the default of 'close = yes' is not viable for an on-line transaction processing environment.
- Only a maximum of 127 rows may be contained on a 4K page; therefore, a row length of less than 32 characters (including overhead) will leave unuseable space in each page.
- Row lengths longer than 2,037 bytes (including overhead) will occupy a full 4K page since rows cannot span pages. DB2 requires 22 bytes of overhead per page; therefore, a row length of 2,038X2 would equal 4,076 bytes plus the 22 bytes of overhead will not fit into the 4,096 page length.
- A row length longer than 4,074 bytes will require a 32K page due to the necessary 22 bytes of overhead per page. 32K pages must be avoided at all costs due to the excessive CPU processing requirements and the amounts of contiguous virtual storage they consume.
- Retrieving a 'column' of data requires CPU instructions, so combining columns in large tables that will have large data retrieval requirements can provide significant reductions of elapsed and CPU times.
- Variable-length columns (rows)
  - Variable-length columns may save large amounts of DASD space.
  - They may also cause higher than expected logging rates for tables that have high volumes of updates. It may not be advantageous to have the variable-length column at the end of the row if other columns are updated with greater frequency. See the next item regarding Logging.
  - They may cause high elapsed times and CPU utilizations for retrievals that select on a set of characters (mask) from the variable-length column or from a column that follows the variable-length element because DB2 must calculate the location of the tested value for *each row* before it can apply the selection criteria.
- Logging
  DB 2 writes all changes (updates, inserts, deletes) to its log datasets. It does not write the entire row when data is updated. It only writes the data from the column updated to the end of the row; therefore, to reduce the amounts of data that will be written to the log when updating, the columns that will be updated with the greatest frequency should be at the end of the row.
- EXPLAIN
  The EXPLAIN option should be used for all queries to evaluate the access path selected by the optimizer. Be sure that RUNSTATS has been executed first so that all statistics are current. PLANS should always be re-bound with the EXPLAIN option. Dynamic SQL access paths should also be evaluated; this may be done through the use of EXPLAIN or by turning on the SQL performance trace and evaluating the generated MINI-PLAN in the Performance Trace Report.
- Indices
  The use of indices must be carefully evaluated and multiple indices should probably not be used for columns that will be used infrequently as selection criteria. An index must be maintained by the system and entries will have to be inserted or deleted every time a row is inserted into or deleted from a table. Using a column having data that will be changed (updated) frequently for an index is simply inviting performance problems — both from the added insert/delete activity in the index and potential page locking in the index. Since an index contains only keys and since there are many more keys per page than would be found in the table, there is also the potential to lock out many data pages in the main table (from users accessing through the index) when locking pages in the index.

Conclusions

We have learned quite a bit about DB2 and many of the design and implementation issues that affect application performance. At the same time, there are questions that arise daily and some of them can only be answered by complicated testing scenarios. A simple change to the structure of a complex SQL request can reduce (or increase) the elapsed time for a query by 50 percent or more.

ABOUT THE AUTHOR

Joel Goldstein is president of The Responsive Systems Company, an information systems consulting firm that specializes in on-line database systems. He has 25 years of data processing experience. Responsive Systems Company, 81 Grand St., Jersey City, NJ 07303, (201) 333-7272.

MAINFRAME JOURNAL
For Users of IBM System/370 Architecture & Compatible Systems

Turn Your Company's Systems Software Tools into New Revenue

We specialize in marketing systems software for IBM mainframes, created by third-party developers. If your company has internally-developed tools and utilities that have marketing potential, we'd like to talk to you.

If your company has developed marketable products in:
- **Storage Management**
- **Programmer Productivity**
- **Data Communications**

call us today for an informal marketing consultation.

Call **Gary Parker**, Product Manager, Sterling Software, Systems Software Marketing Division, at (916) 636-1331, or write 11050 White Rock Road, #100, Rancho Cordova, CA 95670-6095.

STERLING SOFTWARE
Systems Software Marketing Division

MVS Legends *from page 87*

Legends and Rules of Thumb

The legends cited above are qualitative in nature. There are numerous quantitative legends as well. We know them as rules of thumb — usually the really old ones that cannot be explained in current MVS terms.

"There Is a Knee in the Curve"

The Germ of Truth

It is the eternal hope of those who study mathematical functions applying to computer performance that they will look like that in Figure 1.

This might be the case when a process is bimodal and changes rapidly from one mode to another through a small increment of the independent variable. Such might be the response time for a cached DASD volume as cache saturation approaches.

Unfortunately, most functions of interest resemble the M/M/1 queueing function shown in Figure 2.

With a function like M/M/1, there is no critical zone in the domain of the independent variable. The choice of a guideline number is not easy, but the rule-of-thumb makers go right on.

Lesson of the Legend

In most cases there is not a knee no matter how much we wish to find one. Rules of thumb must be questioned if offered without accompanying models that make clear the consequences of violation.

"Keep Device Utilization Below 33 Percent"

The Germ of Truth

If we refer to Figure 2, we see that when the M/M/1 model is an accurate representation of device queueing behavior, a device that is one-third busy will incur queueing delay equal to half its service time. Someone decided many years ago that these numbers had some magical significance — that a device less than one-third busy was not busy enough and that delaying more than half of service time was excessive.

Later Developments and Second Thoughts

Utilization-based performance management methodology has given way to workload-oriented approaches based on classification and ranking of execution and delay states. How busy a device *is* is of less concern than the effect the device has on a key workload's measure of service.

When we do choose to look closely at device behavior, we see many instances in which the M/M/1 assumptions underlying the rule of thumb are not valid. The I/O arrival rate is not exponentially distributed or the service times are more or less deterministic than the exponential model. In every such case, the rule has no validity. A disk drive running 80 percent busy with sequential data transfer is not unhealthy; it is just a funny-looking tape.

Lesson of the Legend

The rule of thumb, applied globally across a large configuration, may be helpful. Certainly it provides a good starting

RD/SHARE® is a complete library management system designed specifically for VM/CMS.

Now owned and marketed by BlueLine Software Inc.

Features include:

- On-line availability of multiple versions of any file.
- Optional full-screen interface.
- Automatic expansion of included files.
- Flexible security structure with a full audit trail.
- JCL wrap facility for interface to guest operation system.
- 100% VM compatible. No CP or CMS modifications required.

Let us share a good thing with you! Call us toll free at **800-826-0313 (612-542-1072 in Minnesota)** for more information on RD/SHARE®.

BlueLine SOFTWARE INC.

1500 South Lilac Drive, Suite 340
Minneapolis, MN 55416
800-826-0313 or 612-542-1072 in Minnesota

point for classifying devices into "over-utilized" and "underutilized" categories. Serious analysis of specific performance problems, however, requires a view of and from workloads rather than of devices alone.

"Keep Path Utilization Below 30 Percent"

The Germ of Truth

What is true for device queueing applies to RPS miss on single paths. The same M/M/1 curve applies with "service time" replaced by "rotational period."

Later Developments and Second Thoughts

Single paths are out, two paths are in and four paths are coming.

Lesson of the Legend

It still applies to MVS/370. In XA, RPS miss requires M/M/c analysis, if not M/G/c. In any event, paths held to 30 or 33 percent busy in XA are often underutilized.

"Keep CPU Utilization Below ?? Percent"

The Germ of Truth

If we look at the CPU delay family of curves in Figure 3, we see another no-knee situation. We thus must pick some "acceptable" level of delay and manage the operating range to stay below that level. CPU utilization guidelines originated with capacity planning considerations when they are useful in establishing the economic order point for a new processor.

Later Developments and Second Thoughts

Percent CPU utilization is of little concern in performance management. The absolute magnitude of CPU delay is of interest and is influenced by the CPU demand per transaction and the contention for the CPU at equal or higher dispatching priority. An added complication comes about when subsystem sub-dispatching is used as in CICS. Such subsystems are vulnerable to internal queueing at relatively low levels of "CPU above."

Lesson of the Legend

Since there is little agreement on the numeric threshold, this legend can probably be dismissed as too vague. CPU delay is usually of little concern, especially in multi-engine configurations. When it

See MVS Legends *page 113*

MHtran-1 VSE to MVS

The most advanced VSE to MVS migration software...anywhere!

For all the details or a FREE demo, call: "The Migration Experts"

MHT Services, Inc.

65 East Route 4, River Edge, NJ 07661
800-MHtran-1; In NJ: 201-342-1321

CIRCLE #46 on Reader Service Card ▲

COMMUNICATE

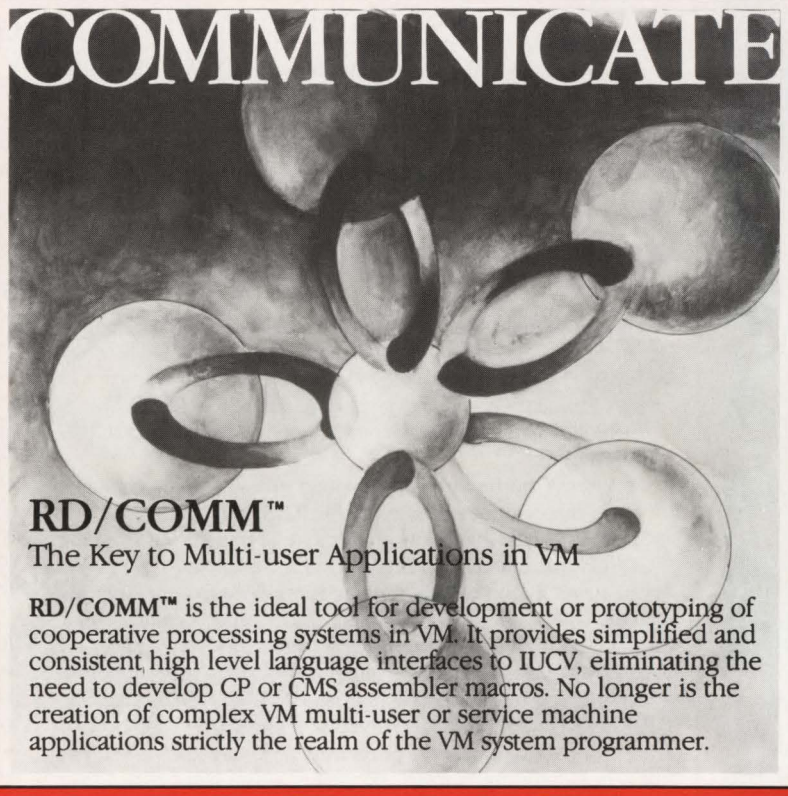

RD/COMM™

The Key to Multi-user Applications in VM

RD/COMM™ is the ideal tool for development or prototyping of cooperative processing systems in VM. It provides simplified and consistent high level language interfaces to IUCV, eliminating the need to develop CP or CMS assembler macros. No longer is the creation of complex VM multi-user or service machine applications strictly the realm of the VM system programmer.

Now owned and marketed by BlueLine Software Inc.

- Distributed with a working sample application, RDUTIL, which provides job scheduler, automatic operator and personal alarm clock services.
- Supports COBOL, PL/I, C, Pascal , and REXX as well as Assembler.
- Requires no CP or CMS modifications.
- Introductory price of $2995.00

Communicate today! Call us toll free at **800-826-0313 (612-542-1072 in Minnesota)** for more information on RD/COMM™.

BlueLine SOFTWARE INC.

1500 South Lilac Drive, Suite 340
Minneapolis, MN 55416
800-826-0313 or 612-542-1072 in Minnesota

CIRCLE #52 on Reader Service Card ▲

TECH ADVISOR

Q *How can I translate a VSAM timestamp that appears in a LISTCAT cluster listing into a readable format?*

A You will need to input the timestamp values into a program to translate them into readable formats. One of DTA's systems consultants, Ed Petschl, has written an Assembler program to translate the VSAM timestamp.

Note to readers: *If you would like a copy of this program, Circle #389 on the Reader Service Card.*

Q *Could you elaborate on STEPLIBS versus LNKLST?*

A The LNKLST is for authorized programs. If you put the authorized program library in the LNKLST, then you can execute the program without the need for a STEPLIB (for that program only).

If you want to run an authorized program without putting the library in the LNKLST, then you must use a STEPLIB for that program and you need to put the library in the APFLIST. But if you put an authorized program in the LNKLST and then (for whatever reason or mistake) you execute the program using a STEPLIB, you lose authorization.

The use of the LNKLST can be a performance issue since the more libraries that are in the list, the longer the search chain becomes for all executable programs.

Q *Can you explain how the report controller feature of CICS 1.7 works?*

A The report controller is an optional feature that can be added (at an extra cost) when purchasing or leasing CICS 1.7 for VSE/SP. This feature provides three major facilities:

1. Additional command-level commands (EXEC CICS SPOOL) to allow the creation of reports from a CICS program.
2. A set of on-line transactions that allows users to manage the reports created by CICS or batch programs.
3. An interface that allows reports to be transferred from the POWER queue to a CICS printer.

The new CICS SPOOL commands are used to create reports that are sent to the POWER queue. These reports can be either LST, PUN or RDR queue entries or transient data. Almost all of the POWER JOB, PUN and LST options are supported when creating these reports; that is, FNO, COPIES, CLASS, DISP and so on. These commands simplify the creation of reports for the application programmer. Using the new functions for submitting to the RDR queue, a user can easily create his own scheduling/job submission software.

The EXEC CICS SPOOL commands simplify the creation of reports in a CICS program and give you the ability to easily specify POWER parameters.

The on-line transaction allows users to display which reports are in the POWER queues, although you cannot view their contents. These reports can then be routed to a CICS or system printer. When printing reports on a CICS printer, the user has the same functions available as on a POWER controlled printer; that is, forms alignments, multiple copies, flush the report and so on.

The on-line transactions work but I find them a bit unwieldy and difficult for users to comprehend. The user functions in the on-line transactions can be limited through security options in the CICS tables. The security is a bit more difficult to specify than it should be. Some third-party print spoolers deal with some of these features more effectively.

The third feature allows a report in the LST queue to be routed to a CICS printer as if it were created by a CICS program. It is useful for short reports in remote locations.

For more information, read the *CICS Release Guide*, GC33-0130.

Q *If a COBOL program uses multiple VSAM KSDS files all using the same SYS number, how do you assign one of the files to 'IGN' from the JCL when using external file names in a DOS/VSE environment?*

A You must change the SYS number on the extent card and add an assign statement assigning the new SYS number to IGN. This will ignore that file. You must be careful using this method of testing since you may inadvertently ignore other files if VSAM also uses the new SYS number you used. A safer method would be to copy the file to a work area and run the job with a different set of DLBL and extent cards. Then you can delete the file when you are done.

Q *We are a VM/VSE shop running a 4381. We have converted to VSE/SP and also converted from BTAM to VTAM in the process. Instead of upgrading our then-current 3274s to SNA devices, we went with SNA 3174s with the assumption that we could accumulate remote response time statistics programmatically. When we asked IBM how to get the statistics, we found that we needed NetView, a $30,000 solution to a $5 problem. Is there an easy (and inexpensive) way to request the statistics and dump them to a flat file?*

A My response is that there are no easy or inexpensive answers to response time statistics. You will pay one way or another for this information.

The least expensive way to get statistics out to a flat file is to use CICS's monitor facility. The manual on customizing CICS explains this very well and it is relatively easy to install. The price you pay is that there are no programs to read this data and condense it into reports. If you belong to any user groups such as GUIDE or SHARE, you can check their cookbooks for available programs; otherwise, you will have to write your own. As you can see, this is not inexpensive but at least there are no out-of-pocket expenses.

The next solution is to buy one of the many CICS response time monitors that are available such as: The Monitor for CICS from Landmark Systems (Vienna, VA), UMAX from Computer Associates (Garden City, NY), Explore/CICS from Goal Systems (Columbus, OH), CICSPARS from IBM, ad infinitum. The advantage here is these products usually have report writers built in with canned reports available.

So far I have only suggested CICS response time options. However the gist of your question seems to indicate that you want network response statistics. I am aware of two options. NetView or a NetView look-a-like product such as Net/Master from Cincom Systems (Cincinnati, OH). Another option would be to buy hardware controllers that have a recording device on them such as Lee Data (Minneapolis, MN) controllers that create a readable file.

All of the above is an answer to a complex question. There are considerations that should be addressed before you start this project:

- What type of statistics do you want? Is it line usage, line problem statistics or task response time?
- Who will want to see these statistics and how should they be summarized — daily, weekly, peak hours?

I would suggest you treat this as a complex project, determine what you want and why you need response time statistics and then establish how much you are willing to spend and determine if the payback is there for your department.

Questions are answered by consultants and instructors from Davis, Thomas & Associates (Minneapolis, MN), the largest technical services firm in the upper midwest. Please address your technical questions to: The Tech Advisor, MAINFRAME JOURNAL, PO Box 551628, Dallas, TX 75355-1628.

CICS *from page 26*

pen. Real storage was increased by 33 percent and processor speed was increased by 40 percent. However it is clear that hardware upgrades alone would *not* have corrected the performance problem. The virtual storage constraint problems in regions was strictly a software limitation. The inability of the storage upgrade to significantly improve response times bears this out.

Exploiting the MRO feature of CICS was the most significant change implemented and it provided a number of benefits in addition to eliminating the virtual storage problem that it was intended to address. Creating the additional address spaces allowed for a higher level of processor utilization without suffering from the software queuing problem identified in Hackenberg's article. The increased number of address spaces also allowed the workload to tolerate a higher paging rate and still provide the sub-second response time required. Simply stated, you get a better return on the hardware investment.

Isolating terminal definitions and dividing the workload also minimizes the impact of a region failure. Failure of an application region will not impact the terminal region or other application regions. In addition, keeping the terminal definitions isolated reduced the time needed to recycle a region. Collectively, these items improve the availability of the workload. Effective exploitation of the MRO feature is a key element in any medium-to-large CICS configuration strategy.

When the processor upgrade did occur, the CICS workload and the rest of the operating system components were positioned to utilize it. Response time improved by 42 percent in response to a 40 percent increase in processor speed. The other components of the operating environment, memory and I/O, were not over committed because of the increased speed of the processor and the multiple regions were able to fully utilize both sides of the 3081 processor. Once again, there was an excellent return on the hardware investment.

ABOUT THE AUTHOR

Frank Bereznay is manager of technical services for the Automobile Club of Southern California. His department is responsible for host/telecommunications system software support, capacity planning, performance management and DASD management.

If you have article ideas, comments or suggestions concerning MAINFRAME JOURNAL, write or call: Bob Thomas, editor-in-chief, MAINFRAME JOURNAL magazine, PO Box 551628, Dallas, TX 75355-1628, (214) 343-3717.

Users Not Complaining...

You'll think it's Magic!

CACHE MAGIC™

A large Northwestern utility company installed Cache Magic, and response times for over 300 users were literally cut in half! And to make sure the improvements were related to Cache Magic, it was uninstalled—the phone never stopped ringing until Cache Magic was back on the job. Now no more phone calls! With Cache Magic, heavily accessed files are stored in processor memory for instantaneous access. I/O bottlenecks disappear!

Why not call today and find out about amazing performance gains for your VM system!

SDI®

20 Years of Software Innovation
1 (800) 537-1969 (415) 572-1200
In Europe: +31 10 436 12 77
1700 South El Camino Real, San Mateo, CA 94402

PRODUCT UPDATE

Powerful Programming Productivity Tool from Integrity Solutions

MAX/SPF, a new programming productivity tool, has been released from Integrity Solutions, Inc. (Littleton, CO). MAX/SPF is a comprehensive, screen-driven, ISPF-based software product that increases the operation and functionality of the TSO/ISPF environment by allowing for faster navigation and easier access to datasets and PDS members.

MAX/SPF addresses the limitations inherent to TSO/ISPF, such as having to remember and enter dataset names over and over, constantly exiting one function and entering another, the inability to edit or browse VSAM files, cumbersome access to PDS files and no effective interactive PDS search capabilities.

For more information
CIRCLE #200 on Reader Service Card

DB2 DASD MANAGER from BMC Software

BMC Software, Inc. (Houston, TX) has announced the availability of DB2 DASD MANAGER, a capacity planning and performance tuning product that manages the life cycle of DB2 physical objects. The latest in BMC's series of DB2 database administration products, DB2 DASD MANAGER provides a comprehensive set of statistics, along with functionality. The functionality is needed to set up objects, monitor growth of tablespaces and execute maintenance when required.

For more information
CIRCLE #201 on Reader Service Card

VMSQL/REORG Announced by VM Software, Inc.

A full-function utility for reorganizing information in SQL/DS was recently announced by VM Software, Inc. (Reston, VA). VMSQL/REORG improves database administrator productivity by reducing the amount of time it takes to perform routine database functions such as; maintenance of databases, DBSPACEs, tables and indexes.

VMSQL/REORG automatically reclusters indexes, reclaims unused DASD, redistributes tables across DBSPACEs, redistributes DBSPACEs across storage pools and updates table definitions. It also features extensive on-line documentation and help files and employs an SAA-compatible, pull-down menu user interface.

For more information
CIRCLE #202 on Reader Service Card

IMS/VS Users Get Integrated DB2 Performance Management

Boole & Babbage (Sunnyvale, CA) has announced IMF (IMS Management Facilities) Release 2.5.0, completing the extensions to all seven IMF products in support of DB2. IMF provides real time, on-line performance management for IMS/VS installations.

DB2 extensions added to IMF Release 2.5.0 are: IMS Workload Analyzer, IMS Workload Monitor, IMS Performance Reporter and IMS Transaction Accountant. These extensions, packaged as options to the base products, now provide total DB2 performance management to installations running both DB2 and IMS systems. The DB2 extensions operate in the MVS/XA and MVS/ESA environments.

For more information
CIRCLE #203 on Reader Service Card

Main Memory Upgrades for 4381 Model 2X Computers

EMC Corporation (Hopkinton, MA) recently introduced main memory products for 4381 Models 21, 22, 23 and 24 computers. These upgrades are available in four and eight megabyte arrays and are 100% compatible with all hardware and software.

EMC's main memory upgrades implement one megabit RAM technology for improved performance and reliability. These main memory upgrades are priced 25% lower than comparable upgrades from IBM.

For more information
CIRCLE #204 on Reader Service Card

PMO/Extended Optimizes Non-ESA DASD Performance

PMO (Program Management Optimizer)/Extended is a management software option from Duquesne Systems (Pittsburgh, PA) that extends processor life and allows non-ESA processors to access DASD as if they had ESA. It reduces end-user response time through improved directory management without the expense of upgrading existing hardware.

PMO/Extended offers performance features for MVS environments such as: directories for all private, non-LNKLST libraries can be kept in virtual storage to further reduce search I/Os; dynamic addition or removal of libraries is accomplished by simple operator commands; on-line listing of candidate libraries; automatic, global LLA refresh.

For more information
CIRCLE #209 on Reader Service Card

Latest Version of EXTEND/VSE Improves Throughput

Improved VSE throughput and increased DASD sharing are benefits of the latest version (2.50) of EXTEND/VSE, a lockfile replacement tool from Goal Systems International Inc. (Columbus, OH).

EXTEND/VSE now reduces I/Os associated with processing VSE Standard/Partition Labels and provides extensive on-line displays that monitor the VSE external lock manager. It also dramatically improves VSE Standard/Partition Label I/O by caching the label area in the EXTEND/VSE server machine storage. DASD contention is reduced by eliminating the physical I/Os needed to handle the request. This capability increases VSE throughput.

For more information
CIRCLE #206 on Reader Service Card

CONTROL-D Report Distribution and Management System

CONTROL-D, an advanced Report Distribution and Management System, recently became available from Tone Software Corp. (Anaheim, CA).

CONTROL-D retrieves reports, using the new CONTROL-D Access Method (CDAM), either from the JES spool queue or directly from the application program. All reports are stored in compressed format. Although some Report Distribution Systems require large dedicated databases for storing and retrieving reports, CONTROL-D requires no preallocated space. Datasets are allocated by CONTROL-D as they are needed and X-37 abends are said to be a thing of the past due to the way CDAM performs dataset allocation.

For more information
CIRCLE #207 on Reader Service Card

CA-APCDOC Automates Production Documentation Requirements for MVS Users

Computer Associates (Garden City, NY) has announced CA-APCDOC, a complete interactive on-line job information system which automatically gathers job, program, procedure and job report information directly from the JCL to provide complete on-line operational information.

CA-APCDOC provides 27 standard reports such as run books, cross references, missing documentation reports, and so on, and includes a flowcharting feature for job scheduling of information stored in the databases of CA's au-

tomated scheduling systems; CA-7, CA-SCHEDULER and CA-MANAGER. It is available for data centers running MVS or MVS/XA.

For more information
CIRCLE #208 on Reader Service Card

New VM/XA Monitor Analysis Program

XAMAP, the VM/XA Monitor Analysis Program, has been announced by Velocity Software, Inc. (Boston, MA). XAMAP analyzes system performance monitor data produced by VM/XA System Product. It generates reports for use in performance analysis activities and also stores, retrieves and reports from extracts of monitor data to facilitate capacity planning and long-term performance trend analysis.

XAMAP reads a file of monitor data collected from a VM/XA SP system. The file may reside on disk or tape. Typically, collection is performed using the MONWRITE command, a tool supplied as part of the VM/XA operating system. XAMAP executes under CMS in either VM/XA or VM/SP.

For more information
CIRCLE #205 on Reader Service Card

VM Users Get the Information Systems Series

Now for VM users, a complete framework for producing custom-tailored standards and procedures is available from Computer Resources Group, Inc. (San Francisco, CA). Called the Information Systems Series, it is an interlinked set of four manuals (and accompanying diskettes) that allows for customization of documentation using word processing software on an IBM (or compatible) PC.

The Information Systems Series helps reduce documentation time, budget and effort because a framework of documentation is provided that has been field-tested and proven effective.

For more information
CIRCLE #215 on Reader Service Card

Westinghouse Enters IMS Software Market

Westinghouse Management Systems Software (Pittsburgh, PA) has introduced two products for the IMS market — SPICE, a productivity aid for implementing restartable IMS/VS application programs and BORIS, a powerful IMS monitoring tool.

SPICE, the first of its kind for the IMS market, is designed to emulate and extend IMS/VS symbolic checkpoint restart and cut the costs of developing and operating restartable applications. It prevents delays in the startup of on-line systems and loss of on-line transactions resulting from re-running of batch work after an initial failure. Jobs can be restarted without JCL changes.

BORIS is an IMS monitor that not only monitors full function and fast path transactions, but also accurately accounts for mixed mode transactions without duplication. It requires no IMS code and imposes no overhead on the on-line system while not impacting the critical IMS logging process in any way.

For more information
CIRCLE #214 on Reader Service Card

Unique Security Product

LOCKSMITH, from SYSMITH, Inc. (Marina del Rey, CA), is a unique computer information security system that controls access to classified information. A technology breakthrough has given LOCKSMITH the capability to recognize, access and protect records, fields and individual characters. This is the key to a series of functions never before possible in security systems because most other systems are not able

See Product Update *page 112*

Batch Runs Done On Time...

You'll think it's Magic!

CACHE MAGIC™

A major Los Angeles manufacturer installed Cache Magic and knocked three hours off their nightly batch processing. They eliminated a two hour daily overrun into online processing, and boosted their system's batch processing capacity by a full hour. As their job mix changed, Cache Magic reacted by uncaching their less active files and replacing them with more active ones — automatically. Instant access to memory resident files meant big performance gains!

Why not call today and find out about amazing performance gains for your VM system!

SDI®

20 Years of Software Innovation
1 (800) 537-1969 (415) 572-1200
In Europe: +31 10 436 12 77
1700 South El Camino Real, San Mateo, CA 94402

SITUATIONS & SOLUTIONS

SITUATION

Improve VSE performance

SOLUTION

If possible, run VSE native instead of under VM. I am not a fan of running a production VSE system under VM. It is simply because of the performance differences.

Migrating from earlier releases of VSE to VSE/SP 2.1 or later will result in a performance improvement for 99 percent of the VSE user base. The reasons for this are: more efficient blocking of VSE libraries, more available virtual storage that should allow most users to increase file buffering, improved channel switching, improved task dispatching at VSE/SP 3.1 VIO POWER queue entries, improved data file blocksizes and a host of other items too numerous to mention.

Improve paging service times or reduce paging?

The first and simplest answer to this question always has to be to add additional real memory. Unfortunately some of us have already reached the VSE limit of 16MB real memory, so for us the solutions are somewhat more complex. If increasing real memory is not an option, then consider the following: improve service times for paging by utilizing faster DASD, eliminate any channel/device contention that may be caused by having other highly accessed data residing on the page dataset device or channel, implement channel switching for the device that contains the page dataset, reallocate the page dataset onto multiple DASD, install a cache control unit for the page subsystem or install one of the memory devices that simulates a disk drive.

If you are running multiple VSEs with any degree of sharing, then a "lock file extension" product can certainly offer substantial performance improvements. When running under VM, there are several lock file products available: CACHE MAGIC (SDI, San Mateo, CA) and EXTEND (Goal Systems, Columbus, OH), and perhaps several others. If you also require a non-VM native VSE system participating in the lock file, then only EXTEND furnishes this support via a CTCA. Utilizing this type of support moves the lock file access to memory, resulting in substantially improved response, a significant reduction in I/O and a consistent leveling of throughput.

Install and implement an improved partition balance facility. Simple VSE time slicing for partition balancing is not conducive to excellent CICS response times. What is needed is a balancer that can be tailored to react to CPU and I/O load.

Resolve the VSE logical transient bottleneck, thereby lessening the contention for resources and allowing more multiprocessing. Several things can be done: load transients into the SVA if room is available, utilize move mode to eliminate the I/O load time and accurately tailor the SDL facility to contain the library addresses of the most-used phases. A better solution is to install the Multiple Logical Transient Area (MLTA) and Fast Transient Loader (FTL) facilities of FAQS from Goal Systems, substantially decreasing LTA contention.

For improved tape processing controls, speed and accuracy, install a tape manager. Upgrade 3420 tape drives to 3480 drives for additional processing speed and improved reliability. On 3480 tape drives there is no real space or speed advantage unless blocksize exceeds 4K. Performance and speed continue to improve with larger blocksizes. The minimum recommended is 16K; preferred, 32K.

To improve label processing efficiency and reduce label processing I/O and waits, implement a facility to place the standard label areas in memory. Many installations do not realize how large a bottleneck this can be. If you cannot place the label area in memory, then ensure that labels are located as rapidly as possible by placing them into the user label area. Your worst option is to place all labels in the standard label area since it is always searched last. Searching the label area is done via a search cylinder command with a key equal option and can result in the DASD being busy for a lengthly interval. Label area search order is user, partition and standard. Utilize user everywhere possible. Do not place the label area on a DASD that already contains other highly accessed files or facilities.

Install a terminal for the operator that can support an alternate console facility. The primary reason for this is so the operator can utilize the alternate console for scrolling back and forth while the system console can continue to write messages and jobs can continue to process rather than wait because a message needs to be output and the console has been scrolled backward. There are several excellent VSE third-party vendor console products on today's market. Of course, this could be accomplished utilizing the console support VSE/SP contains in the interactive User Interface. Many of the console products offer special user defined PF key support at the native console. Implementing this for the 24 most-used commands will certainly speed up operator responses.

Another solution is to implement a facility that would answer most, if not all, console replies automatically. Try ZACK from Altai Systems (Arlington, TX). Another possibility would be to implement a batch control file where answers could be stored prior to batch runs and have the applications access the batch control file rather than the console for replies.

Significant improvements in batch run times can usually be obtained by reducing all required interaction between the system and an operator. Several users have asked if they could use the IBM supplied Job Manager facility.

Yes, you can use it but be aware that IBM does not support your use of this facility. It is supported only for IBM's specific use. The Job Manager is not documented and is limited to only one active partition at a time. I believe it is of limited value as a job manager facility in a production environment. If you really need

a job manager/scheduler, there are several good ones on the market. The one I am most familiar with is another Altai Systems product called ZEKE. ZEKE is an exceptionally powerful job scheduling system that I believe can manage even the most complex environment.

Another possibility for a rather simplistic job manager system is to utilize VSE/SP conditional JCL, the VSE job accounting interface, and a user-coded batch program as processing controls. I have seen some rather simple inventive "Job Manager Systems" created using these facilities.

Generally a VSE user should attempt to minimize I/O that must be performed to service a request. A VSE user should also minimize any contention that would result from having highly accessed resources together. Always attempt to separate POWER queue and data files, the page dataset, label areas, VSE system libraries, the VSE lock file and highly accessed data files.

SITUATION

Improve POWER performance

SOLUTION

The POWER data blocksize should be enlarged to more efficiently transfer POWER data. I recommend using 4K or 6K as a good balance of memory utilization versus DASD space utilization for users of VSE/SP 2.1 or later. For prior releases, just under 2K seems best.

POWER users prior to VSE/SP 3.1 should investigate BIMPDQ from B I Moyle & Associates, Inc. (Minneapolis, MN). BIMPDQ puts the queue file into virtual memory for POWER 2.2 and earlier users and dramatically reduces I/O request. POWER 2.3 implements a similar facility using the V/O facility of VSE/SP 3.1. Both facilities offer substantial performance improvements.

If you cannot utilize one of the in-memory queue file facilities, then *do not* put the POWER data and queue file on the same DASD spindle and *do not* put either one with any other highly accessed data. The POWER files tend to be one of the top three-to-four busiest packs.

Implement one of the "fast printer support" facilities. The one I am familiar with is contained in FAQS from Goal Systems. FPS blocks data coming from a partition to POWER, thereby reducing interruptions to POWER for spool-out requests.

To reduce POWER processing overhead, keep the active queue file entries to a minimum; long queue chains require substantial resources to process. Do not use POWER as a long term storage facility for printouts; offload them to tape.

SITUATION

VSE Enhancement Task Force

COMMENTS

Membership on the Task Force depends on developing, testing, implementing and distributing a successful VSE extension. If you are interested or have suggestions, let me hear from you (615) 622-5141.

ABOUT THE AUTHOR

Pete Clark has been in data processing for 25 years, the last 11 with Olan Mills, Chattanooga, TN. Clark is a recognized authority on the DOS/VSE operating system. He has worked many years to extend the limits of VSE.

Performance So Amazing...

You'll think it's Magic!

CACHE MAGIC™

With Cache Magic, improved VM system performance has never been easier. Enjoy accelerated online response times, dramatic reductions in batch run times, and complete more transactions in the same, or less time—all without the cost of additional hardware. Cache Magic eliminates I/O bottlenecks by storing your most active files in processor memory. And Cache Magic is always at work, automatically determining which files should be cached.

Why not spend ten minutes with one of our representatives to find out what Cache Magic can do for you? All you have to lose are user complaints, batch overruns, and expensive hardware upgrades.

SDI®

20 Years of Software Innovation

1 (800) 537-1969 (415) 572-1200

In Europe: +31 10 436 12 77

1700 South El Camino Real, San Mateo, CA 94402

CIRCLE #76 on Reader Service Card ▲

to detect any increment of data smaller than a file.

LOCKSMITH is a low-level system, completely separate and insulated from applications code. It permits a much faster and easier implementation by not impacting either the program code, the database structure or the program architecture.

For more information
CIRCLE #212 on Reader Service Card

MAGEC Application Development System Enhanced

MAGEC Software's (Plano, TX) Mask & Application Generator & Environment Controller (MAGEC) software, which is driven by its active data dictionary, now contains extended dictionary facilities which provide automatic field-level help for all MAGEC generated on-line applications. MAGEC automatically generates basic help text which may be supplemented with additional explanations by the user, if desired.

This new version (2.0) also includes a new feature for the database administrator called Business Rule Processing. This new facility enables the DBA to include unlimited data validation and default rules along with the basic data definitions. To remove the burden from the application programmer, the processing is automatically inserted into any generated applications which update that data element. These benefits are available using standard VSAM, with or without a third-party DBMS.

For more information
CIRCLE #213 on Reader Service Card

DASD/Plus Offers Continuous CICS Availability

DASD/Plus is a CICS resource management system that lets data center operators carry out traditional CICS facility management tasks while providing application users with continuous system availability. Recently introduced by On-Line Software International, Inc. (Fort Lee, NJ), DASD/Plus also significantly reduces the twin liabilities of off-line maintenance and application down-time traditionally associated with CICS management.

DASD/Plus provides real-time support for files and transient data queues, which must be defined to the system in order for CICS to access them. It also performs on-line file allocation and deallocation.

For more information
CIRCLE #211 on Reader Service Card

JCLPREP Announced by Altare

JCLPREP, a complete JCL management system, has been announced by Altare Computer Company (Gilroy, CA). JCLPREP detects JCL errors, reformats JCL, automates JCL conversions, dynamically modifies JCL and provides validation and enforcement of standards. It comes complete with generic reports and auditing and has built-in report writing capabilities.

At the heart of JCLPREP is the unique Expert Editing Facility (XEF) which will perform logical JCL manipulation and modification. It requires no compiling, assembling, linking or programming expertise.

For more information
CIRCLE #210 on Reader Service Card

WHAT ZEKE DOES FOR AUTOMATED SCHEDULING...

When Zeke, the scheduler that works, was released in 1982, it revolutionized automated job scheduling. And today, it's still the best scheduler on the market – better than CA-Scheduler, UCC-7, or any other.

Zeke is designed for today's data centers, not yesterday's batch environments. Only Zeke can open and close online files. Automatically reply to job messages. And bring the online system up and down.

With Zeke, installation takes hours instead of days. Implementation takes days instead of months. Whether you're running MVS, VSE, or both.

And there's another important difference: Zeke costs less. Less to buy, less to implement.

For a free copy of Datapro's independent analysis of job scheduling software, call Altai Software today.

You'll find out how Zeke revolutionized automated job scheduling – and how it can do the same for your data center.

ZEKE

The scheduler that works.

CIRCLE #114 on Reader Service Card ▲

VMOPERATOR *from page 102*

"Since then, we have taken advantage of many other features," Columbia says. One result is that Pfizer no longer keeps hardcopy operator logs: it is all on disk. It keeps about a month's worth of data on disk just as it did when it had hardcopy. VMOPERATOR overwrites the oldest data automatically. Columbia is gratified that at the top of every screen, VMOPERATOR lists the CPU utilization and the number of users logged on to the system.

Pfizer has configured VMOPERATOR so the operator can control any disconnected machine. By hitting a PF key, operators become the operators of disconnected machines such as RSCS or VTAM just as if they had logged on directly. "The beauty of it is that it will not let you log that machine down from the VMOPERATOR environment," Columbia says.

Columbia is also pleased with the Names feature of the newest release. This feature displays all user IDs on the system sorted by CPU utilization, Start I/Os and pages. "Now, when the resource utilization is high or response time is low, we can use VMOPERATOR as a monitoring system to see what the problem is," he says. For more information, contact VM Software, Inc., 1800 Alexander Bell Drive, Reston, VA 22091, (703) 264-8000.

ABOUT THE AUTHOR

John Kador is a free-lance writer and a frequent contributor to MAINFRAME JOURNAL.

MVS Legends *from page 105*

does turn out to be significant, workload-oriented approaches should lead to useful results quickly.

Legends of Tomorrow

As we wait to see what MVS/ESA will bring, we note that the announcements were short on specifics especially about performance. The opportunities for legend formation abound.

"MVS/ESA Requires Two Gigabytes of Expanded Storage"

Beginning of a Legend

This assertion appeared in a trade press article based on a few interviews several days after the announcement. What someone said or what the reporter concluded was that because MVS/ESA can exploit a great deal of expanded storage (the two gigabytes available on the largest 3090Es), all of that is required.

Lesson of the Legend

The assertion is false and dangerous. As facts replace speculation and unwarranted inference, it should fade away.

Conclusion

Legends are necessary in history and civilization. They serve a limited purpose in MVS performance management. When they obscure facts and rule out viable choices, when the underlying facts have changed, legends must be discarded or at least rebuilt from scratch. If the advice given throughout this article can be summarized, it is "Listen, learn, analyze, plan, measure — and think for yourself!"

ABOUT THE AUTHOR

Stephen L. Samson is an executive consultant for Candle Corporation, 1999 Bundy Dr., Los Angeles, CA 90025, (213) 442-4150.

© Copyright Candle Corporation 1988.

ZACK WILL DO FOR AUTOMATED OPERATIONS.

Introducing Zack, the operator's operator—a revolutionary new product from Altai Software.

Zack lets you literally *program* your computer for unattended operations.

Other products may duplicate a few of Zack's capabilities—but none come close to matching Zack's combination of outstanding features and exceptional flexibility.

Zack not only suppresses and responds to messages. Zack also lets you issue system commands. Write your own operator commands. Schedule commands by time of day. And more. Which means greater productivity, efficiency, and accuracy for any MVS or VSE system.

Working alone, Zack can achieve dramatic results for your data center. And for even greater benefits, you can use Zack side-by-side with Zeke, the scheduler that works.

To find out more, call Altai Software today. We'll show you how the operator's operator can automate *your* operations.

ZACK

The operator's operator.

ALTAI

Software that works.

ALTAI Software · 624 Six Flags Drive · Suite 150 · Arlington, Texas 76011
1-800-227-7774 (answered 24 hours a day)

CIRCLE #150 on Reader Service Card ▲

VENDOR PROFILE

SDI

Twenty Years of Software Innovation

Twenty years in the life of most industries is merely a point on a time line; a fraction of its overall history. Not so in the world of computer software. Here, twenty years represents the entire life of the industry. And a dynamic life it has been; one marked with rapid, sometimes fatal changes.

It is for these very reasons that SDI is proud to celebrate its twentieth anniversary this year in the business of providing high quality, innovative systems software to the IBM mainframe market. In 1968, SDI was among the handful of companies to pioneer the then unique concept of selling software to mainframe users. Founded by three former IBM professionals, these men understood the need for software products that enhanced and expanded on those provided by IBM. They also understood that users were willing to pay for software that could offer hardware cost savings and improved productivity.

SDI develops and markets DOS and VM software products designed to fill a specific "niche" in the market: products that provide immediate, benefit-oriented solutions to user dilemmas. That is a philosophy that is not going to change in the near future, at least according to SDI's Bruce Pastorius, director of marketing.

"Our approach is consistent and relatively simple. All of our products need to meet four basic criteria. First, they must be innovative. There are no "me too" products in our current line. We listen to our users and directly address their problem areas. In fact, some of our products have no competition within the industry," says Pastorius.

Secondly, our products must be simple to install by our customers. We take great pride in ease-of-use and implementation. Thirdly, they have to be easy to maintain — no complicated procedures to cloud product utilization. Lastly, and most importantly, every product must supply readily quantifiable benefits to each and every user," explains Pastorius.

SDI has maintained an enviable track record with a list of products that have kept pace with IBM hardware and operating system developments. Their first product, GRASP, a comprehensive family of programs designed to enhance performance and hardware utilization in DOS environments, was eminently popular with IBM System/360 users. Many well-received products followed including EPAT, the world's first automated tape library manager. As the mainframe user migrated to the System/370 and DOS/VSE, SDI introduced Instant-FBA, the first software product designed to provide complete disk compatibility for DOS/VSE users.

More recently, SDI took notice of the fact that VM was becoming the operating system of choice for many mainframe users and began to address some of their unique problems. SDI entered the VM systems software market in 1984 with VM Magic, a disk compatibility product. VM Magic allows any guest operating system under VM to utilize any physical DASD without regard for previous restrictions caused by disk architecture, access methods or programs. VM Magic's popularity was confirmed by an ICP Million-In-One Award in 1985.

The newest addition to the SDI product line is Cache Magic. Cache Magic enables VM users to dynamically place their most heavily used disk files in processor storage for instantaneous access. This enables them to greatly reduce I/O operations and enhance overall system performance without costly hardware upgrades.

"We fully realize our products are somewhat esoteric and that educating our prospects in their capabilities is part of our job. Still, when potential users find out they can run any operating system with any disk device or that they can improve overall VM system performance overnight, their skepticism disappears," Pastorius points out.

From the beginning, one of SDI's priorities has been to provide product compatibility with IBM's new and extended operating systems. This commitment was reaffirmed recently when SDI announced support for its VM products in the VM/XA environment.

SDI realizes its products must be technically excellent to complement and enhance IBM's systems software. The products are conceived and produced exclusively for IBM or compatible mainframe users and all interface with standard IBM software. It is corporate policy to maintain the simplest of installation procedures and strive for the most expedient and solution-oriented technical support possible. Each product is supplied on a standard format magnetic tape accompanied by a clear, easy-to-understand User Guide. Support is available 24-hours-per-day, 365-days-a-year.

The scope of SDI's product line is constantly expanding. "It takes a great deal of research and analysis to transform a product 'possibility' into a functioning, solid software product. We know that well-designed systems can mean the difference between productive performance and costly bottlenecks. Good ideas and customer feedback are key, but ideas alone do not get the job done — it takes people," says Pastorius.

"I would say we are unique in the industry in that we have individuals in marketing, development, technical support and administration who have been with the company for more than ten years. When you add that experience level to a track record for breaking new ground, you have a powerful combination — stability for our customers and a legacy of innovation for our own people," Pastorius reasons.

SDI-USA is headquartered in San Mateo, CA with field sales offices in Hackensack, NJ and Minneapolis, MN. In addition, SDI has an international division based in Europe and has agents in Canada, South America, Australia, Africa and Asia. The company has manufacturing facilities in California and the United Kingdom with research and development centers devoted exclusively to the design and manufacture of systems software products.

By listening to its customers, SDI feels able to respond to their needs in the most timely means possible. "We take pride in our ability to pinpoint operating system shortcomings and in the ability of our development team to create products to fill these gaps. We are specialists, if you will, in extending and improving the capacity and performance of our users' mainframes. We have been at it for twenty years and we are proud of our accomplishments. We hope to be doing the same twenty years from now," concludes Pastorius.

Vendor Profile is a regular forum whereby a vendor is given the opportunity to introduce the company and its products to MAINFRAME JOURNAL readers.

STROBE

MVS APPLICATION TUNING

Exploring the Boundaries of Space-Time

Basic questions about Space-Time remain unanswered today. Will the Universe expand forever? What, if anything, could retard its expansion?

What about your IBM MVS mainframe universe? Will your memory, disk space, hiperspace and CPU time requirements also expand forever? Is there anything you can do to retard this expansion?

Fortunately, you don't have to master Einstein's general theory of relativity to answer these questions. Programart's STROBE® Performance Measurement System tells you just where CPU time is spent within your mainframe address spaces. It enables you to quickly spot and correct inefficiencies that produce excessive resource demands.

With STROBE, you can:

- ☐ Tune your applications both during and after development
- ☐ Reduce online response times
- ☐ Shorten critical batch processing runs
- ☐ Postpone costly hardware upgrades

To learn how to keep critical applications within your space-time boundaries, call Programart today at (617) 661-3020.

PROGRAMART

1280 Massachusetts Ave. ■ Cambridge, MA 02138

STROBE will help you tune applications in CICS, DB2, IMS DB/DC, and other DB/DC environments, as well as batch processing programs.

STROBE is a registered trademark of Computer System Architects Incorporated.

© Computer System Architects Incorporated 1988.

CIRCLE #33 on Reader Service Card ▲

It's time we shed a little light on what we're up to.

For years you've known VM Software as "The VM Experts"—the #1 source of data center management software for VM environments. Today we're more committed than ever to the needs of VM users. And we're exploring a whole new range of ways to help—both inside the data center, and outside.

Some of these new directions have even taken us beyond VM itself, to the complex world of multi-operating-system networked environments. As these environments become more common, the need for new kinds of systems management tools becomes critical. And we are now providing the first of these tools, with many more to come.

What lies ahead? A continuing stream of innovation across four product areas with crucial implications for VM and networked environments: Data center management. Relational technology. Network data movement. And network management and administration.

We believe these directions are in sync with your own evolving needs. As those needs continue to evolve, rest assured that we will be there with products and technology you can count on. For more information write or call VM Software, Inc., 1800 Alexander Bell Drive, Reston, VA 22091, (703) 264-8000.

CIRCLE #20 on Reader Service Card ▲