MAINFRAME JOURNAL

For Users of IBM System/370 Architecture & Compatible Systems

August 1989

Using DB2 At Glaxo Inc.

UNLEASH THE POTENTIAL OF DB2 AND SQL/DS.

A full set of database tools to enhance performance and simplify administration.

Making the most of your DBMS is a lot easier with the right tools. Now there's a full set available from Systems Center for two of today's most popular relational environments: DB2 and SQL/DS.

Our DB2 software products (DB/SECURE,™ DB/AUDITOR,™ DB/REPORTER,™ and DB/OPTIMIZER™) address urgent needs with innovative, effective solutions — streamlining security, simplifying auditing, speeding report generation, and boosting system performance. All while eliminating the errors and delays associated with manual DB2 administration.

In the world of SQL/DS, our DB/REORGANIZER™ product dramatically enhances performance and efficiency in a fast-changing database environment. Our DB/EDITOR™ offers exceptionally convenient full-function table editing. And our DB/REPORTER,™ with its outstanding data manipulation and formatting facilities, makes even complex reports a relatively simple matter.

So why wait? Start making the most of your environment — and yourself. Call or write today: Systems Center, Inc., 1800 Alexander Bell Drive, Reston, Virginia 22091.

800-359-5559 703-264-8000

SYSTEMS CENTER

A NEW YORK STOCK EXCHANGE COMPANY

RELATIONAL DATABASE PRODUCTS VM SOFTWARE PRODUCTS NETWORK DATAMOVER PRODUCTS NETWORK ADMINISTRATION PRODUCTS

DB/AUDITOR,™ DB/EDITOR,™ DB/OPTIMIZER,™ DB/REORGANIZER,™ DB/REPORTER,™ and DB/SECURE™ are trademarks of Systems Center, Inc. and its subsidiaries.

CIRCLE #20 on Reader Service Card ▲

6-MFJ-8908

You need a whole MVS system, not one that's full of holes.

The New Monitor for MVS.

Without a complete monitor, you're vulnerable to holes in your management of MVS/XA and ESA. Holes that can lead to nasty falls.

Now, with The Monitor for MVS™, you can easily fill those treacherous gaps.

TMON/MVS gives you the equivalent of five other monitoring products, with more extensive coverage than you can get anywhere else.

You get an exception monitor, a top-notch delay monitor, and a DASD monitor to see real-time events in your system.

Real-time information isn't enough. So TMON/MVS gives you something unique — an online, systemwide record of the recent past. That's backed up by a performance data base for online query and batch reporting.

Whether you're running a few monitors or none, if you're not running TMON/MVS, you may have gaps you can't see. Avoid the pitfalls, with TMON/MVS.

For a FREE TRIAL call 1-800-227-8911 or (in Virginia and Canada) 1-703-893-9046.

THE MONITOR FOR MVS™

The Whole Solution

LANDMARK™

Landmark Systems Corporation
8000 Towers Crescent Drive
Vienna, Virginia 22182-2700

CIRCLE #16 on Reader Service Card ▲

CONTENTS

August 1989 • Volume IV • Number 8

DEPARTMENTS

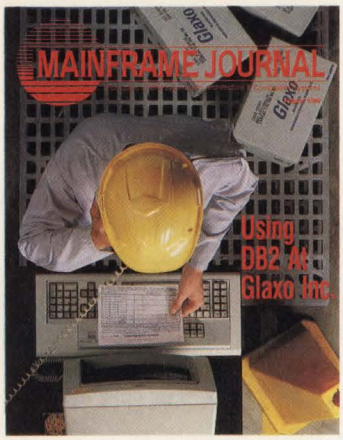

COVER:

Glaxo Inc. researches, develops and manufactures prescription medicines. Its successful strategy to implement DB2 is outlined starting on page 19. Cover photograph provided by Dawn Winter, Glaxo Inc.

MAINFRAME JOURNAL© (ISSN 0895-5751) is published monthly by Thomas Publications, Inc., 10935 Estate Lane, Suite 375, Dallas, TX 75238, (214) 343-3717. Second class postage paid at Dallas, TX. SUBSCRIPTION RATE: Subscriptions are free within the USA and Canada. One-year foreign subscriptions are $96. POSTMASTER: Send address changes to: *MAINFRAME JOURNAL*, P.O. Box 551628, Dallas, TX 75355-1628.

MANAGEMENT

SYSTEMS

APPLICATIONS

OPERATIONS

With CA-LIBRARIAN You Have As Much Change Control As He Does. Almost.

Our library management system has evolved to a higher plane of existence. CA-LIBRARIAN® is now the only comprehensive software package that combines library management and change control for all IBM operating environments.

You can now rely on CA-LIBRARIAN to manage the thousands of program changes taking place in your source library every year. You get on-line audit capability, automatic synchronization to ensure every change is valid within your system. Built-in security reinforcement. Plus, a complete record of code changes made on a day-to-day basis.

To date, more than 7,000 facilities have come to rely on CA-LIBRARIAN. Find out why by calling Jamie Brooks at 800-237-9273 (in N.J., 201-874-9000). And discover what it's like to have total control over your destiny.

Well, almost.

COMPUTER® ASSOCIATES
Software superior by design.

PUBLISHER

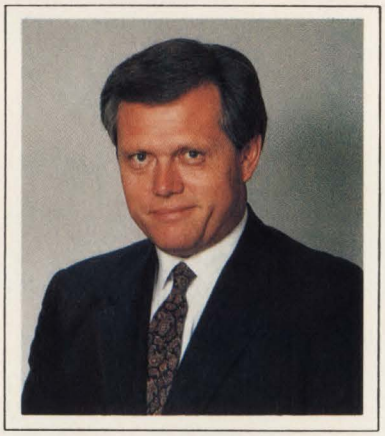

Bob Thomas

What's An IDUG?

IDUG — It sounds like the monosyllabic utterance of some ''hip'' jazz buff nostalgically remembering a great Charlie ''Bird'' Parker performance in the '50s, preparing to scrawl ''BIRD LIVES!'' across a New York City facade.

Actually, IDUG is the obligatory acronym of the International DB2 Users Group. IDUG is assembling in Chicago this month (August 6-9) for its first-ever conference, so we decided to provide a DB2 emphasis in this issue of *MAINFRAME JOURNAL*.

Robin Pasley, DBA for Glaxo Inc., details the way Glaxo is using DB2 on page 19. Robin is quick to point out that ''not all of the applications are currently executing on a relational platform and some may never be.'' In Barry Lewis' article on page 38, ''DB2 Storage Groups,'' he says that the combination of DB2 Version 2's new features and the use of partitioned tablespaces may allow the storage group to become a viable production alternative. Product Review on page 76 focuses on how BMC Software's DB2 MASTERMIND can help manage the complexity of the DB2 environment. Our final DB2-related article features the current President of IDUG, Bill Backs, Director of Information Technology of Scott & Foresman, who steps up on the *MAINFRAME JOURNAL* ''soapbox'' (that is, Viewpoint, page 98) and attempts to debunk a couple of myths concerning DB2/relational DBMSes.

VSE Hits The Big Time

Over the past five or six months, it has been interesting to observe and participate in the VSE-user ruckus. To put it succinctly, many die-hard VSE users were feeling left out by IBM and ''they weren't going to take it any more!'' The charge was led by several individuals, most visibly Pete Clark, Charlie Rice and Eric Vaughan — all past or present officers of the VSE contingent in GUIDE. Pete went so far as to actually provide the code necessary to extend the lagging capabilities of VSE when it became doubtful if IBM would.

As the furor built, ''lowly VSE'' became front-page news in such prestigious publications as *Computerworld* and *Software Magazine*. Modesty prevents me from detailing the numerous VSE-related articles run in *MAINFRAME JOURNAL* and the tremendous number of copies of ''Pete's code'' sent out gratis to all who asked.

What was IBM doing while VSE users ''burned?'' Well, apparently they weren't just fiddling around. They were listening! Even IBM can't afford to have 24,000-plus VSE users starting to ''circle the wagons.''

We were recently invited to White Plains, NY for a refreshingly candid briefing on IBM's position on VSE. Their designated spokesman, Bernd Robatzek, Director of VSE Development in Boeblingen, West Germany, was forthright in admitting a miscalculation of the discontent felt by VSE users. (See an interview with Robatzek on page 12.) He went on to state that VSE is strategic to IBM's future. It is viewed as important not only as a stand-alone product, but also as a key component, along with the 9370, of larger mainframe-based networks loaded with PS/2 workstations to fuel their growth into the '90s.

To paraphrase our jazz buff friend — ''VSE LIVES!''

Bob Thomas

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

Publisher and Editor-in-Chief
Bob Thomas

Associate Publisher
Martha Thomas

Editor
Carol M. Hoag

Copy Editors
Judy Beller
Pat Warner

Art Director
David Kramer

Assistant Art Director
Ken Buerer

Magazine Production
Diane Dishman

Marketing Services
Sally Webb

Circulation Manager
Janice Porter

Assistant Circulation Manager
Nancy Crawford

Administrative Manager
Marian Davenport

Advertising Sales
Denise Thomas — Mark Cauto

BPA

List Rental Information
Contact AllMedia (214) 985-4060

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.

The Power Package!

Take the shortcut to fast response times. DATAPACKER/II, the CICS transmission optimizer from H&M, will reduce outbound and inbound data streams by 50-85%. The advantages speak for themselves: considerably improved response times, no screen blinking, lower line costs and free capacity for additional terminals and new applications.

DATAPACKER/II
CICS · TRANSMISSION · OPTIMIZER

CIRCLE #118 on Reader Service Card

H&M
SOFTWARE'S FUTURE

H&M Systems Software, Inc., 25 E. Spring Valley Ave., Maywood, N.J. 07607-2150, 1-201-845-3357, 1-800-FOR-DEMO

Lights-Out Operation Debated

I found Howard Miller's question-and-answer article on unattended computer center operation in the April 1989 issue interesting and informative. One theme that runs through the article, however, bothers me: the ''absence of people'' in the data center.

I would suggest that unattended operations means that *peripheral* operations are unattended. Leaving the system running, but unattended, is a mode of operation that I would not advocate. In addition, I don't think it is a good business practice. It would seem that this method of operation implies no on-line systems, therefore, no users present to require services.

My plan for a lights-out/unattended operation would entail the following:

• Convert all tape datasets to DASD and eliminate the tape drives. Install an STK-type device in a remote location to handle the DASD backups and archived data. Before you dismiss this as too costly, have you ever analyzed tape data and volumes to see what is really needed?

• Convert all paper/microfiche reports to on-line viewing. This will reduce the requirement for high-speed laser printers and report distributions or pickup. The remaining printing, check *stubs,* A/P checks (if required) and so on will be printed on *department* printers.

An important consideration in this area is that most applications should be redesigned to accommodate on-line viewing. Merely shifting 132 print positions (or more) from paper/microfiche to a CRT, in many cases, is not the best way. Remember, in many instances we are dealing with *batch systems* that have been around for quite awhile. In addition, more CRTs will likely be required, since, in many instances, those who use terminals today are either doing data collection (entry) or are viewing other files.

Many installations have a data control function where ''run-to-run'' totals, balancing and so on is carried out. This function should be automated through software by Systems Development or through a software package. What remains in the data center is a system control function and a network control function. In many installations job scheduling, data center scheduling (holidays, power shutdown, testing) and so on is handled within the data center; in some others an administrative group (under Data Center control) provides these functions but is located outside the data center. This group also interfaces with the users. At this point, it may be most practical to locate all remaining data center activity *outside* the data center. Then you can really be ''lights out.'' If these functions remain in the data center, it is still lights out (everywhere but the console area).

I feel that ''lights out'' or ''unattended operations'' refers to peripheral functions only and was never intended to reduce service. Obviously, a satisfactory level of service must be provided regardless of the mode of operation. Any change of operation that users perceive as a reduction in service levels (that is, no one around to provide help) must be addressed to prevent dissatisfaction.

The idea behind unattended/lights-out operation is to increase performance and quality, since manual functions and human intervention are minimal.

Roger Claus
Hallmark Cards Inc.
Kansas City, MO

VM Bigotry

As a certified VM Bigot, I am always glad to see VM articles in *MAINFRAME JOURNAL*. I was disappointed to see that the examples used in Michael Seadle's article, ''VM In The Development Center,'' in the June 1989 issue used a poor coding technique by depending on the REXX default of setting the value of a variable to the name of the variable if it is referenced before it has been set to a value. To use one of his examples, take the MICHAEL exec from the center of page 96. The next-to-last line reads:

```
EXEC NOTE MICHAEL '(' NOTEBOOK PROJECT
```

This line contains five REXX variables (EXEC, NOTE, MICHAEL, NOTEBOOK, and PROJECT) and one literal string ('('). Of these variables, only one (PROJECT) has been set to a value. The other four will default to their literal value. A better way to code this line would be:

```
'EXEC NOTE MICHAEL (NOTEBOOK' PROJECT
```

The line now contains one literal string and one variable (PROJECT) that has been set to a value. As shown in the first example, the exec will work properly and for such a simple exec it didn't really matter. It's when you have more complicated execs running to several hundred or even thousands of lines that strange things can happen. Suppose you have a large exec that Joe from another department wrote awhile ago with a line like the first example buried within it. Now you are changing it to add a new option and you need a variable; so you insert the line:

```
NOTE = FRITZ
```

Now you have introduced a bug that is going to be somewhat obscure and difficult to track down because when you get down to the portion that wants to send the note, the exec will try to execute an exec named ''FRITZ'' that may or may not exist on your system. Even worse, if ''FRITZ'' had previously been used as a variable name, the problem is compounded. Also, if the exec loops, it might work the first time around the loop and fail in strange ways the second time.

To sum up, if an item in the REXX exec is a literal value, make it a literal instead of a variable by putting it inside quotes. Either single quotes (') or double quotes ('') can be used, but make sure you use the same type at each end.

Rich Greenberg
The Traveling VM SYSPROG
Los Angeles, CA

Editor's note: In response to Greenberg's suggestions, Seadle says it boils down to a question of style. Greenberg's way is preferable for those who lose track of their variables; Seadle's way allows you to give variables new meaning if desired.

KUDO

This was my first issue of *MAINFRAME JOURNAL*. I'm very pleased. As luck would have it, we are currently working on change control. I found your article, ''San Bernardino County's Checklist For Change Control,'' (June 1989) especially useful.

F. L. Langley
Computer Task Group
Kirkland, WA

New 3270 & Token Ring Support

Remote printing on a PC without a single compromise.

Until now you've had to rely on a S/36 or AS/400® to deliver remote printing. Now a PC with BARR software and adapter sustains print speeds of 6,000 lines-per-minute and line speeds of 128,000 bits-per-second. Only BARR maintains all this performance with the reliability and ease of use PC users expect.

BARR/SNA RJE or BARR/HASP software drives up to five printers from a single PC. What's more, you can enter data, print, and receive output all simultaneously – without interruption. BARR's advanced multi-tasking software easily manages even the most complicated tasks, including LAN access, tape support, file transfer, and special forms printing. In addition, BARR offers one year of friendly, dedicated customer support with each purchase.

Communications adapters and software are available for the IBM PC, PS/2, and compatibles.

Try BARR for 30 days. We've helped thousands save millions. Call 800-BARR-SYS.

Protocols		
SNA RJE with Multiple Sessions		
HASP – BSC Multileaving		
Host Systems		
JES2	DOS/POWER	CDC NOS
JES3	RSCS	VS1/RES

BARR

BARR SYSTEMS INC.
2830 NW 41 Street, Gainesville, FL 32606
800-BARR-SYS, 904-371-3050, FAX: 904-371-3018

AS/400 is a trademark of IBM Corp.

CIRCLE #91 on Reader Service Card ▲

Computer-Related Security Risks Increased

According to Ernst & Whinney's recently completed 1989 Computer Security Survey, 71 percent of the respondents reported that computer-related security risks increased in their organizations over the past five years. This is a 21 percent increase over last year's survey. Half of the respondents reported financial losses over the past two years because of system failures or downtime and about a quarter of them reported financial losses because of malicious acts. Regarding security technologies, more than 82 percent of the respondents use optional access control software such as RACF and ACF2. Seventy-three percent use Uninterruptible Power Sources (UPS) and almost 63 percent employ redundant communications or power to avoid single points of failure. Surprisingly, only a few respondents reported plans to use some of the more recent (and expensive) biometric security technologies such as fingerprint and retina scanners. Twenty-eight percent say they plan to use virus detection software. These findings suggest that organizations are attempting to find non-expensive, "quick fix" solutions to difficult computer security problems. *Source: Ernst & Whinney, Cleveland, OH.*

Open Systems: The Greatest Industry Force in the Past 25 Years

The information technology industry has beheld numerous trends, fads, so-whats, sure-things and wholly unexpecteds, but none of the developments have the long-term consequences of the trend toward open information systems. According to Peter L. Burris, editor of *The Gray Sheet Computer Industry Report,* "Open systems is the greatest industry force witnessed in (our) 25-year history." He defines a system to be open if characterized by two conditions: (1) it is unable to exert a controlling influence on its operating environment and (2) facilitating technology transfers is its fundamental organizing principle. Much open systems' activity in the market is geared towards Unix, not so much by design as by fate. Although its high portability and cheap licensing costs made it a natural choice for many hardware startups, the industry is not heading toward open systems because of Unix per se. The real acid test for determining Unix's place as customers strive to build more efficient systems is: can the open systems' movement deliver Unix mainframe functionality through an open process? *The Gray Sheet* thinks not. Developing, building and supporting mainframes is a highly integrative activity, probably beyond the scope of the existing open systems movement. Besides, the S/370 architecture clearly dominates the mainframe arena enough to be called standard and IBM's Systems Application Architecture (SAA) certainly has many attributes of openness. *Source: The Gray Sheet Computer Industry Report, International Data Corporation, Framingham, MA.*

Disaster Recovery Services: A Market With Potential

Only about 15 percent of the IBM/PCM mainframe shops in the USA and Canada are currently using an outside disaster recovery service while six percent of the sites analyzed are planning the use of such a service in the next 12 months. The popularity of such services varies by industry with the banking industry lead followed by the transportation/utilities and the insurance/finance industries. Currently, only 13 percent of sites with a single system and 19 percent of the sites with multiple systems have opted for the assistance of disaster recovery services in protecting their data processing operations. This reveals the possibility of a large untapped market for companies offering this type of service. This potential might explain why IBM has chosen to enter this market as well. *Source: Computer Intelligence, La Jolla, CA.*

XA Systems Corporation and KPMG Peat Marwick Announce A Strategic Alliance

XA Systems Corporation, a provider of productivity software tools for IBM's MVS environment and KPMG Peat Marwick, the world's largest accounting and consulting firm, announced the formation of a strategic alliance. The alliance brings together XA Systems' XPERT Series of testing and maintenance productivity software and the PMAT CASE and Reengineering product line. Under the terms of the alliance, XA Systems will assume full ownership of the PMAT CASE and Reengineering products and have complete responsibility for the marketing, sales, support, distribution and development of the product line. KMPG member firms around the world will provide associated consulting assistance in conjunction with the implementation of the expanded XA product line.

Computer Associates Security Products Support IBM's DFP Release 3.1

Computer Associates International, Inc. recently announced that its CA-ACF2 and CA-TOP SECRET access control products are fully functional and support IBM's Data Facility System Managed Storage (DFSMS)/DFP Release 3.1, a key component of the MVS/ESA operating system. This announcement coincides with the general availability of the enabling PTF for DFSMS functionality in a DFP 3.1 environment.

DOCS™

THE BEST-KEPT SECRET IN VSE CONSOLE AUTOMATION.

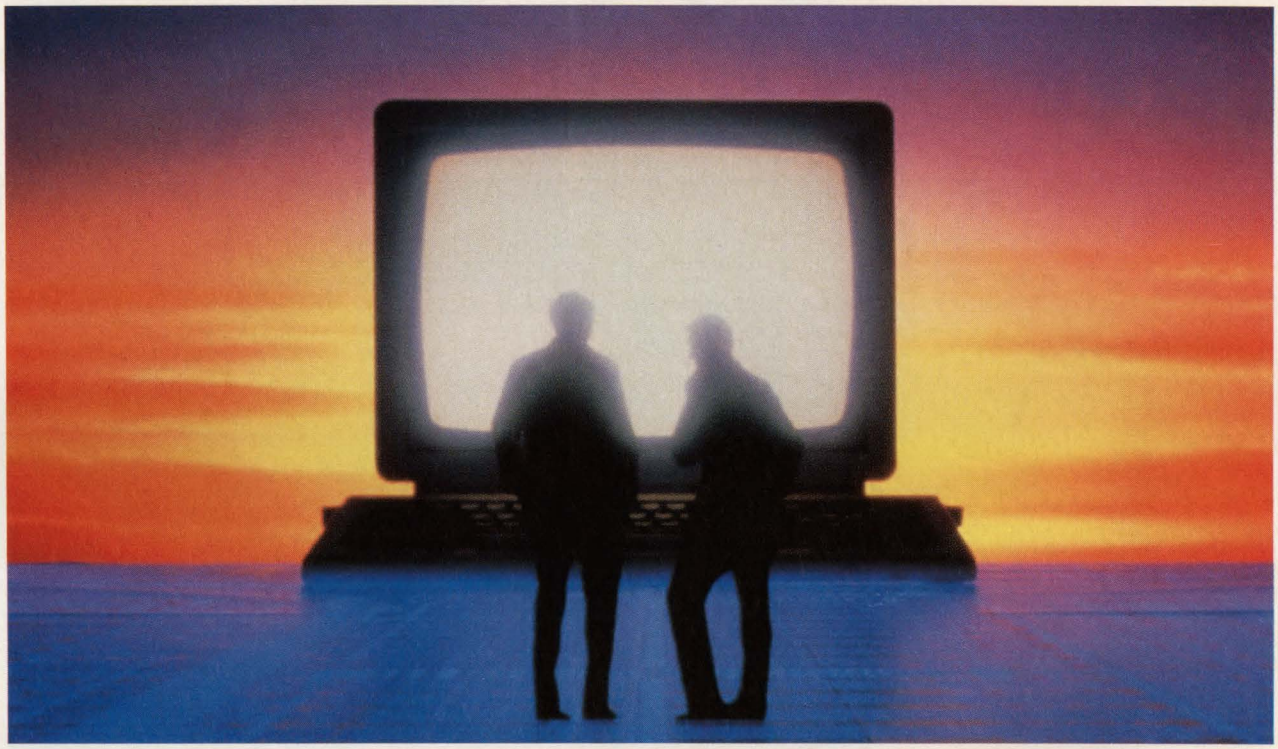

Today's data centers are serious about managing their VSE operations. Most realize the system console is the logical first step. But what many don't realize is that the solution for VSE console automation is already in use at hundreds of data centers around the world.

SINGLE IMAGE CONSOLE

DOCS from SMARTECH SYSTEMS features a Single Image Console Facility under VM (SICF/VM) that lets you control multiple VSE systems from one DOCS console. Multiple VSE consoles are consolidated on one CRT, creating a more efficient operation. *And you always have access to local or remote consoles* because DOCS is not dependent upon an online system.

The VM OPERATOR console can be consolidated and managed from the DOCS/VSE console, which means you view VM operations in full-screen 3270 mode using one less CRT. You can redisplay VM operator messages online and list VM hardcopy data because DOCS logs *all* messages to the hardcopy file with date and time stamp. You even have access to CMS minidisks and can execute REXX programs directly from the VSE console.

MESSAGE SUPPRESSION AND ROUTING

DOCS message suppression and routing capabilities let you customize your console to display only the messages you need to see. You can suppress messages at selected terminals, or route messages based on partition, device or user profile.

...AND MORE

DOCS offers more than *fifty additional features*, making it the most comprehensive VSE console automation solution available, Features such as auto-reply, programmable function keys, security features, unattended operations capabilities and more.

NO RISK, MONEY-BACK GUARANTEE

Buying DOCS is risk free because it comes with an unconditional six-month money-back guarantee. *And all it takes is a simple 30-minute installation.*

So if you are looking for the secret to VSE console automation, find out what hundreds of our customers already know. For more information on your console automation needs, call 1-800-53-SMART. (Outside the U.S., call 214/956-8324.)

SMARTECH SYSTEMS, INC.

Turning high technology into *SMART TECH*nology.℠

10015 W. Technology Blvd., Dallas, TX 75220
FAX: 214/357-6338 Telex: 9102503110

International Representatives:

France: Futurs Systemes et Technologies • Poitiers, France
Tel: 33-49-01-7374, FAX 33-49-01-7351

Australasia: Mycroft Systems, Ltd. • Auckland, New Zealand
Tel: 64-9-817-7673, FAX: 64-9-817-3640

United Kingdom: Rendeck U.K. Ltd. • Essex, England
Tel: 44-0702-333555, FAX: 44-0702-333055

For additional international representative information, contact SMARTECH Systems, Inc.

IBM Speaks Out On VSE

An Interview With Bernd Robatzek

Bernd Robatzek, Director of VSE Development for IBM in Boeblingen, West Germany, agreed to answer questions posed by MAINFRAME JOURNAL. The questions stem from a concern on the part of VSE users that IBM's direction and focus were shifting away from VSE. They were concerned because it appeared that VSE was not being enhanced, thus forcing many users to migrate to MVS prematurely. According to Robatzek, IBM will listen more intently to VSE customers on an ongoing basis. VSE is strategic and is being enhanced.

IPL Introduces 8MB Memory for 4381 Processors

At IPL, we've built a 15-year reputation of providing low-cost, high-dependability IBM compatibles. We not only produce fully-compatible memory cards for all your 4381 models; we also guarantee them with a lifetime warranty and overnight replacement.

Our 4381 boards are available in 8, 16 *and* 32MB 4-card sets, giving you maximum expansion flexibility.

Call IPL today for specification and ordering information, and join other 4381 users who already count on IPL for reliable, low-cost memory cards.

Call IPL today at 1-800-338-8ipl, in MA (617) 890-6620. Or contact us at 360 Second Avenue, Waltham, MA 02154.

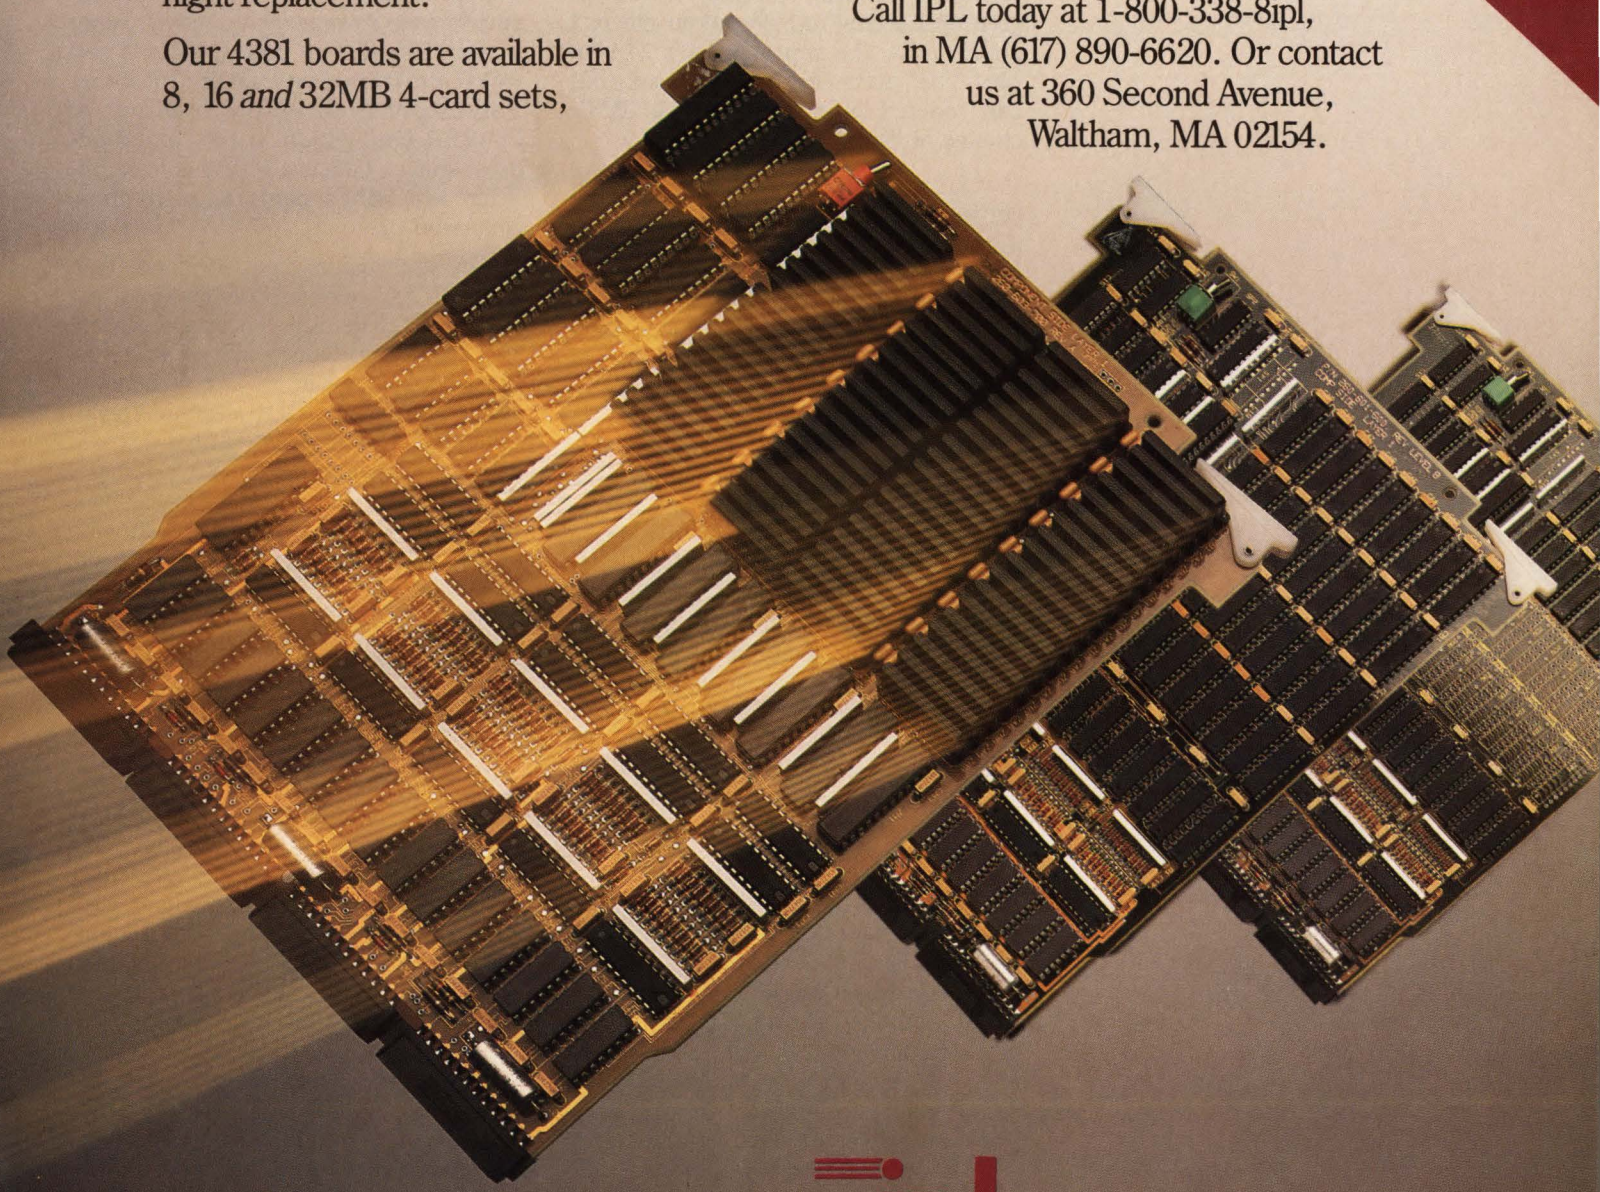

ipl systems, inc.

IBM is a registered trademark of International Business Machines Corporation.

CIRCLE #96 on Reader Service Card ▲

Q A

We will further improve the affinity between MVS and VSE operating systems by concentrating on subsystem compatibility and common user interfaces.

Q Many VSE users felt abandoned when it appeared obvious that VSE was not being enhanced nor was it included in products related to IBM's "strategic direction;" however, recently a 180-degree turn seems to have taken place. I am sure that the approximately 24,000-plus VSE users are quite pleased, but can you tell us what prompted IBM's change of heart?

A VSE has always been a strategic product for IBM and our customers and at no time did we plan not to continue to support our VSE customers. For a time, however, the importance and the growth potential for the transaction-oriented environment of VSE was in question and this had reduced the emphasis on VSE extensions.

However, we have been listening closely to our VSE customers about what value VSE can offer on solutions for the future. Among the things that have prompted a stronger focus on VSE is, for example, the growing importance of a transaction/database-oriented control system as a versatile server in the fast-growing world of interactive, intelligent workstations.

Q VSE users would like for there to be an affinity between the MVS and VSE operating systems; what plans does IBM have for this?

A Today, VSE already offers good affinity and compatibility with MVS, particularly in the CICS transaction-processing environment. Major subsystems like CICS, COBOL, SQL, CSP, VTAM and so on are supported in VSE as well as in MVS. We have published a brochure that describes how to develop compatible applications between MVS and VSE.

VSE, besides its traditional role as a S/370 intermediate data center, is more and more used as a decentralized, distributed MVS system.

We will further improve that affinity by concentrating on subsystem compatibility and common user interfaces. However, there will always be more functional richness on MVS.

Q At a recent GUIDE session, many VSE users strongly expressed needs concerning larger partitions (more than 16MB virtual) and more partitions as well; in addition, there is a need for support for real storage above 16MB. What is your response to these concerns and what time frames are realistic? Another VSE-user concern is the residence of POWER and VTAM in the shared address space. This is a severe constraint. What are IBM's plans to solve these concerns?

A The user requirement for more than 16MB of real storage is on the top of my priority list. The user concerns regarding larger partitions can be solved with various approaches and we are in ongoing discussions about how to stage the implementation. Let me mention that this is a topic we also discuss with our customers. We just had a worldwide VSE customer council in Boeblingen, West Germany, where our planning and development organization shared ideas and discussed priorities.

Q Based upon the "new direction," what will be VSE's relationship to SAA? Will VSE ever be a full participant in SAA?

A In September 1988 we formally announced the relationship of VSE to SAA. This announcement describes the SAA participation of VSE for the VSE/CICS environment. For the SAA Common User Access (CUA), the Common Programming Interfaces (CPI) and the common communication interfaces, this means the following.

For the CUA interface, VSE will be using and integrating the programmable workstation. For the CPI, we have announced the SAA subsystems COBOL, CSP and SQL. Furthermore, we have also announced all communication interfaces for VSE. Similar to MVS, we also made a statement of direction to provide support for a programming workstation interface for CICS applications (CPI-C).

As SAA evolves, we will also be evaluating improved relationship and participation of VSE in SAA.

SYSTEMS SOFTWARE FOR VM/VSE DATA CENTERS:

Only one company covers the two, completely.

Computer Associates presents: **CA-UNICENTER®/II VM** and **CA-UNICENTER/II VSE**, the industry's only complete line of VM/VSE systems software that automates every major area of the data center.

Now, true compatibility among the VM and VSE components. Equivalent products for both environments offer the VM/VSE data center unparalleled advantages such as: a common catalog that simplifies tape management, security software that protects all data in your installation, job accounting information that is collected for activity in both VM and VSE, and much more.

Only Computer Associates provides common interfaces and full integration to give you unprecedented control, from a central point, over both environments.

And only Computer Associates offers CA-UNISERVICE®/II, a secure link between your mainframe and CA's Customer Service System, 24 hours a day. You get online access to software fixes, interactive problem resolution, product tutorials and more. No one else has anything like it.

Call Dana Williams today:
800-645-3003

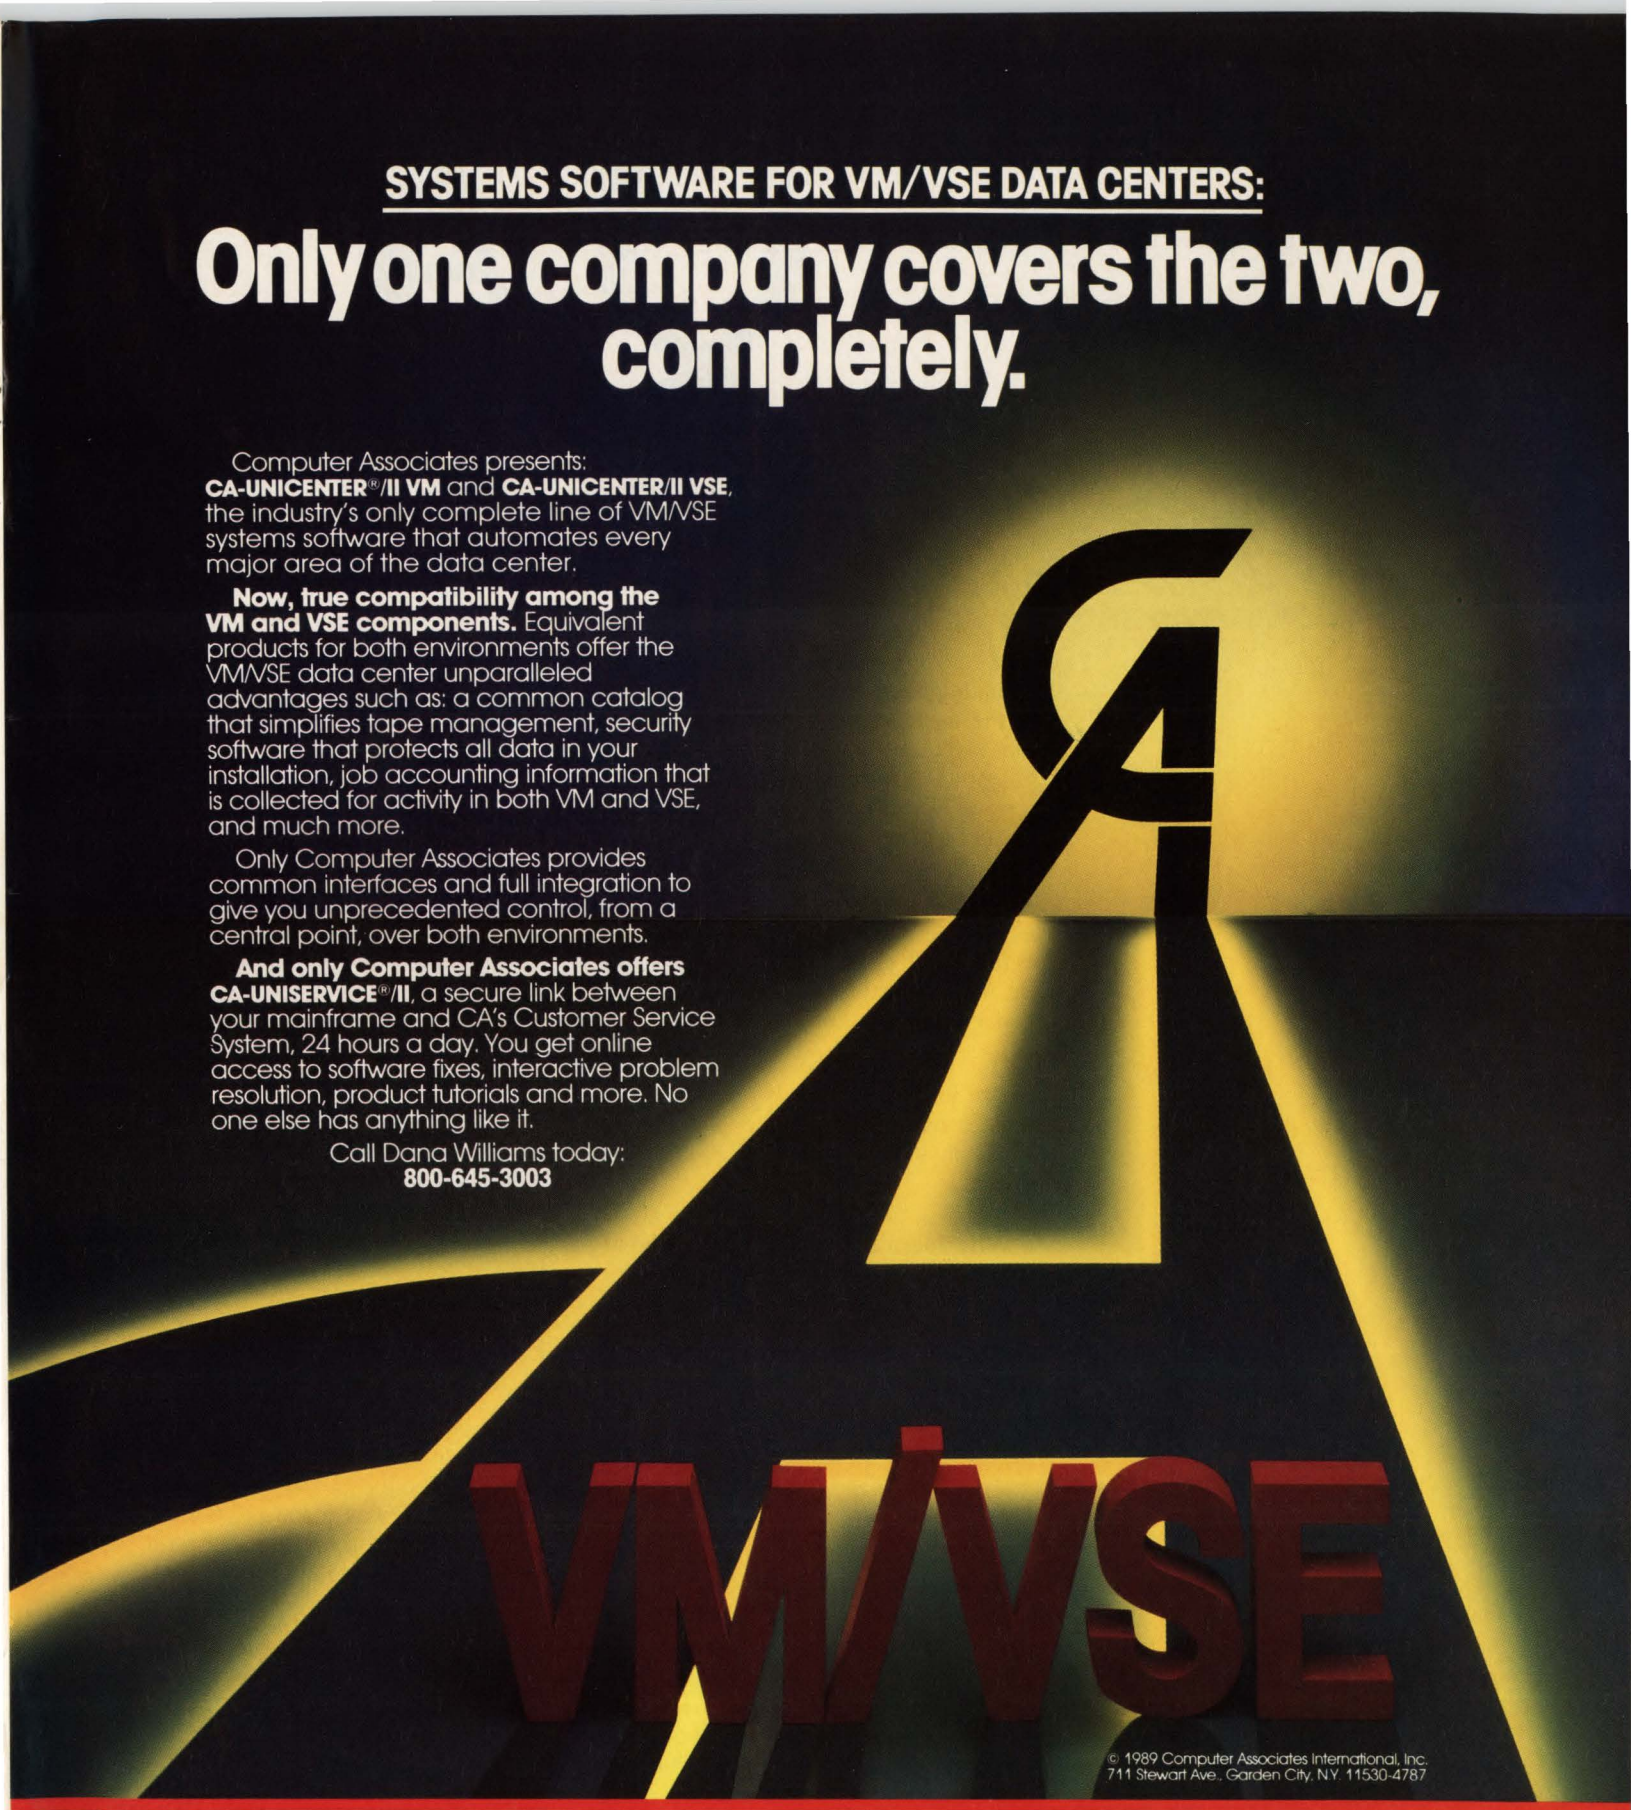

COMPUTER ASSOCIATES®
Software superior by design.

- 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 #90 on Reader Service Card ▲

Q A

As SAA evolves, we will also be evaluating improved relationship and participation of VSE in SAA.

Q Will VSE ever support ESA?

A Together with the user requirements listed under question three, we are also evaluating ESA support for VSE. I do not see the need for all ESA functions to be made available in VSE. If we made all ESA or MVS functions available in VSE, then we would build a second MVS system and I believe we already have a very good one.

If we made all ESA or MVS functions available in VSE, then we would build a second MVS system and I believe we already have a very good one.

Q What is IBM doing to improve the performance of the VSE DASD subsystem with regard to the limitation of the number of paths and channels?

A This is an accepted requirement. The level of implementation is to be evaluated under consideration of the architectural implementation of future S/370 midrange systems.

Q Why doesn't the new OfficeVision offering support VSE?

A We just started to ship a modular VSE office package ("VSE Office Offering") for non-programmable workstations this year. This provides good functions for the user with existing 327X workstations and is compatible with the current MVS office products. The new OfficeVision includes the capability to connect the VSE office function with OfficeVision via local-area networks. Additional office enhancements for VSE are under evaluation.

Q VSE/SP 4.0 is due out momentarily and 4.1.1 is scheduled for December. What additional features do you envision being added over the next 18-month time frame?

A Our first priority is to support our installed customer base with new hardware support, constraint relief, enhanced function for ease of use, distributed VSE systems, comprehensive networking and easier migration to MVS.

Also of high priority is the support of key SAA elements, improved MVS affinity and improved application solutions by working closely with software vendors and system integrators.

Q From your perspective, what do you see as VSE's future with regard to growth? Do you see a significant number of new VSE users and where will they come from?

A The VSE customer base has been steadily growing over the years. Recent growth comes mainly from the ES/9370 system where VSE is used as a distributed system. For the future I see new VSE users in the area of transaction application solutions, programmable workstation integration and, last but not least, server concepts.

Q Many VSE users have complained to us and in GUIDE sessions about the lack of interest and concern IBM support and service people seem to exhibit toward VSE accounts. The implicit and explicit reference to VSE being "non-strategic" and the lack of timely receipt of VSE-related announcements evidence this lack of concern. What is IBM's response to this situation?

A I think we have proven a strong strategic focus for VSE by making eight major VSE announcements during the last 14 months. This was made possible by listening to and working together with customers and user organizations.

Q Please comment on the large MVS users' need for the VSE operating system, specifically as it relates to the distributed environment.

A The announced support for distributed and unattended systems with VSE Version 4, the inclusion of VSE NetView and other MVS communication products and the portability of CICS applications all provide an attractive distributed offering in areas where the functional richness of MVS is not required.

Why ADABAS 5 is thousands of dollars faster.

The investments your organization makes in its data base technology will either cost it a fortune—or save it one. That's why you need a DBMS that assures optimum performance in a high production environment: ADABAS 5.

In a recent series of standard, fully scaled benchmarks, audited by Coopers & Lybrand, ADABAS 5 proved again that it is thousands of dollars faster. Each benchmark was conducted on a National Advanced Systems AS/EX™100 (equivalent in power to an IBM 3090 500S).

In the **standardized TP1 debit-credit benchmark,** ADABAS 5 worked with a data base containing **1 million accounts.** The results: a **record-breaking 388 transactions per second** (tps)—99.3% serviced in under one second!

For the first time, an authentic **ET1 debit-credit benchmark** was conducted with ADABAS 5 managing **10 million accounts.** The results: **167 tps** (from terminal in/to terminal out)—99.5% serviced in under one second!

These figures represent major savings in time and money. But they're not surprising—at least not to the thousands of organizations which already use ADABAS 5.

They expect more performance. Plus the benefits that come from 18 years' experience in DBMS technology:—relational functionality which allows adaptable data structures and fast information retrieval based on multikey criteria—document management and free text retrieval—24 hour processing in a fully integrated DB/DC environment—portable applications across various hardware and operating systems—distributed processing—entity relationship data bases with recursive retrieval functions for knowledge-based systems.

Discover how much more profitable your DP organization can be. **Conduct your own ADABAS 5 benchmark,** using your own data and application profile in your own production environment. The facts will speak for themselves.

™ AS/EX is a trademark of National Advanced Systems Corporation.

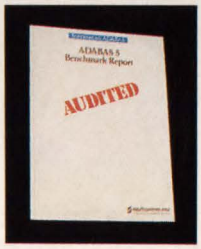

Demand the performance and functionality only ADABAS 5 can offer.

To order your free copy of the ADABAS 5 Benchmark Report, call toll-free:
1-800-843-9534 today
(In Virginia or Canada, call 703-860-5050.)

SOFTWARE AG
PROGRAMMING BUSINESS SUCCESS

CIRCLE #16 on Reader Service Card ▲

Unlock The Mysteries Of DB2 Performance.

To manage DB2 performance, you need more than information. You need solutions. OMEGAMON® for DB2 gives you both.

Exception Analysis warns you when key thresholds are exceeded. Before a stray thread turns your system to stone. Then powerful zooming features take you where no one else can. Deep into the passageways of DB2. Right down to the SQL call. And Recommendation Screens translate detailed data into realtime solutions.

Yet what's sophisticated on the inside is made simple on the outside with menus and help screens. OMEGAMON identifies the problem thread so your troubleshooting is as effortless as pressing a PF key.

Proprietary software engineering keeps overhead as low as 1% with minimal space requirements. And OMEGAMON is always available—even when DB2 is locked up.

Candle's always available, too. With round-the-clock customer service, technical education, and a commitment to stay current with IBM. So you won't ever be lost in the passageways of DB2.

To decipher the mystery of DB2 performance, call Terry Forbes today at (800) 541-8513.

Candle®

CIRCLE #116 on Reader Service Card ▲

Using DB2 At Glaxo

By Robin Pasley

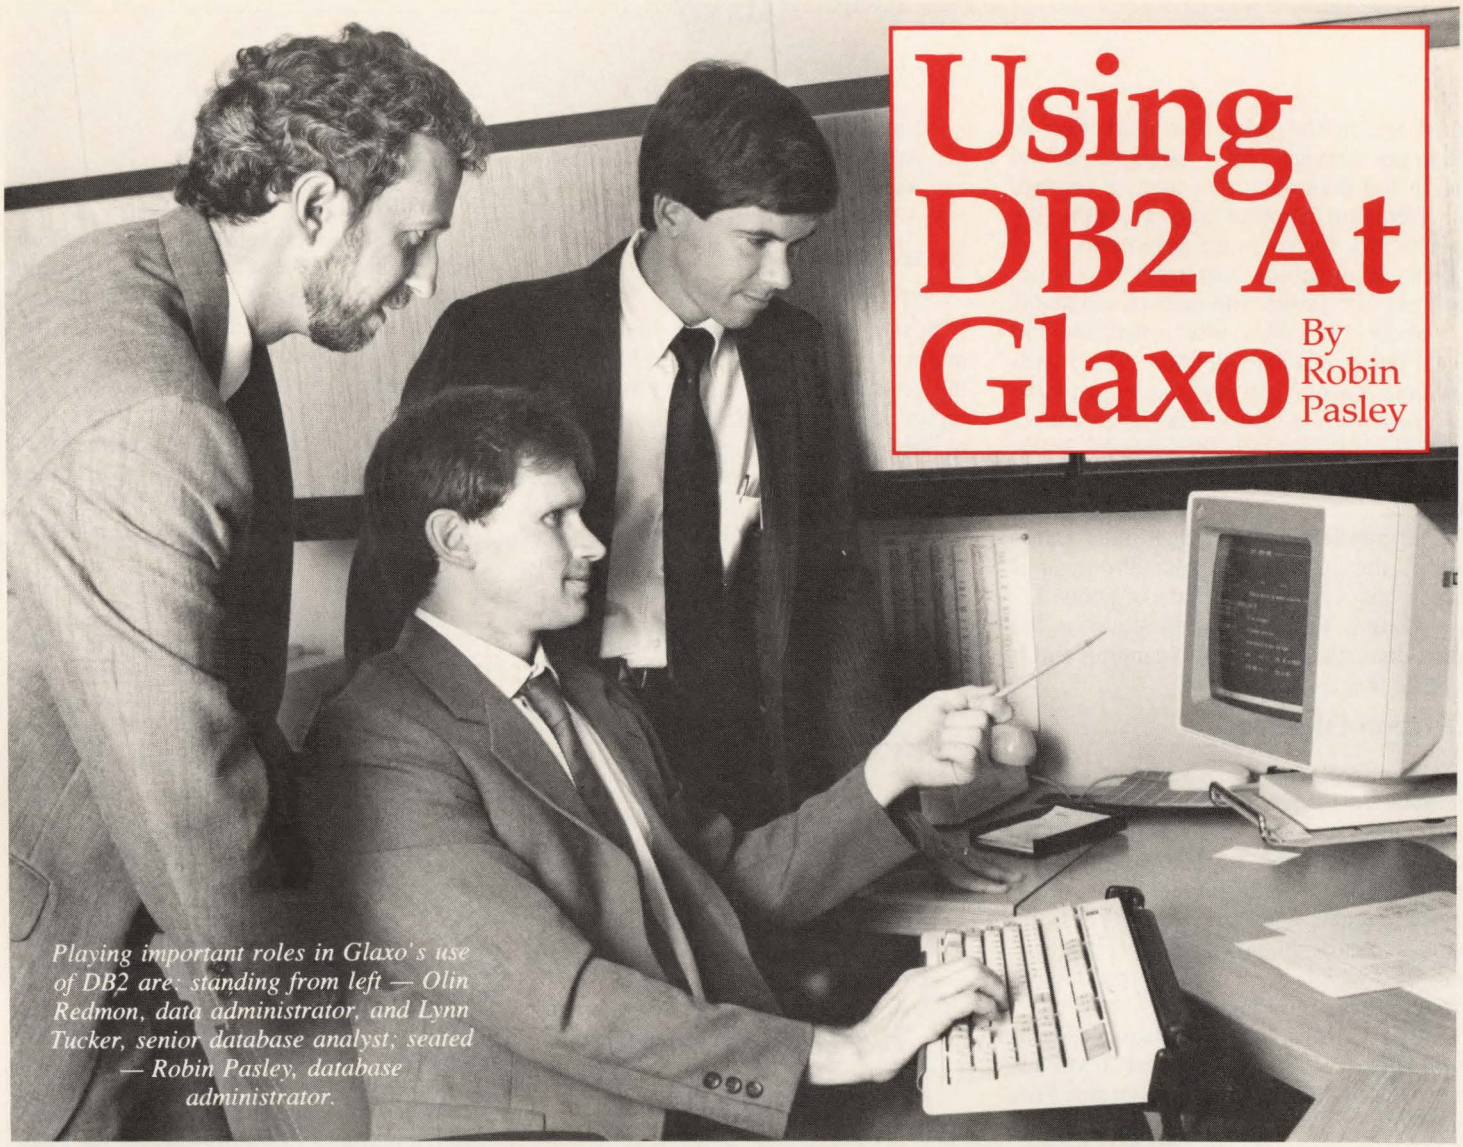

Playing important roles in Glaxo's use of DB2 are: standing from left — Olin Redmon, data administrator, and Lynn Tucker, senior database analyst; seated — Robin Pasley, database administrator.

Glaxo Inc., with headquarters in Research Triangle Park, NC, researches, develops and manufactures prescription medicines, including those for the treatment of gastrointestinal disorders, respiratory ailments, cardiovascular diseases, infectious diseases and diseases of the skin. It is a subsidiary of Glaxo Holdings p.l.c., of London.

Glaxo has standardized on SQL-based relational technology. DB2 is used on an IBM mainframe for all new commercial application development and Oracle on a VAX platform for most new scientific and research application development. This article details the way Glaxo is using DB2.

There are three points worth making before I begin the meat of the article. The first one is that not all of the applications are currently executing on a relational platform and some may never be. An application that is running fine, already paid for and has little or no need to exchange information with other applications will probably not be rewritten until replacement is required for some reason besides the fact that it does not use a relational DBMS.

Secondly, the nature of Glaxo's business is such that it does not have the same high-demand-level transaction processing requirements as an airline or a telephone company. The most demanding transaction-based application currently planned for DB2 will have a transaction to process every 1.2 seconds.

Last of all, vanilla DB2 as provided by IBM does not include facilities to manage the product as effectively as Glaxo would like. Several after-market tools were incorporated into the environment to help overcome perceived shortcomings. The tools will be discussed as the problems they help resolve are presented.

For the DB2 environment, the measure of success is based on various standards. One is easy availability to all data without artificial ''application'' boundaries with security facilities to prevent unwanted access. Another standard is the ability to cleanly divide labor and responsibilities along lines that promote maximum productivity and accountability.

A third one is support for a mixed environment consisting of both transaction processing and decision support applications sharing the same data whenever practical. This is desirable because of the time and resources consumed in building and maintaining ''bridges'' to other software and/or other data structures.

Acceptable performance is the last standard. *Acceptable* means different things to different people. The three major categories of interest are transaction processing, decision support and batch processing. Each category will be addressed independently later.

Glaxo runs DB2 on an IBM 3090 using MVS/XA. At this time, the company is using Version 1.3 of DB2. There are 28 DB2 applications in production status.

Most are in the sales and marketing or customer service areas; however, a few are in the manufacturing area, a few are for Decision Support and a couple are specifically for software developers.

The applications currently in production use 305 tables; the ones in test status use about 325 tables. Also, there are 350 or so personal tables defined. The smallest table has one row and the largest table has two million rows. The largest table currently fills four double-density 3380s and is expected to fill eight by the end of the year.

The application requiring most transaction throughput is Order Processing, which has a transaction rate of about 1.2 per second. The application requiring the most disk space is Image Scanning and it will be discussed later.

Division Of Labor

The effort required to support DB2 is divided into units providing strong support and accurate accountability with minimal headcount. Glaxo is, in fact, trying a somewhat different, less centralized approach with Oracle and scientific applications development. Experience will indicate which is the more productive strategy over the long haul.

However, this article deals exclusively with Glaxo's DB2 environment. I will address the traditional areas of Systems Programming, Computer Operations, Data Administration, Database Administration, Security Administration, Applications Development and Quality Assurance, explaining how each area relates to DB2 support at Glaxo.

Glaxo is divided into three discrete groups which cover all areas mentioned above. The first is Systems Services that includes the Systems Programming and Computer Operations functions. Next is Information Resource Management that includes Data Administration, Database Administration and Security Administration. The third group is Applications Development.

We have no staff dedicated strictly to Quality Assurance. The following Quality Assurance functions are distributed among the three groups above as follows:

- Application code reviews are done by Applications Development
- Migration to production is done by Operations and Database Administration
- Data quality assurance is the responsibility of Data Administration
- Naming standards are enforced by Applications Development, Data Administration or Database Administration as appropriate.

System Services

Systems Programming

Systems Programming installs and maintains DB2 and many related tools. All products which require special training or authorization to install are handled by this group. This includes products from IBM and anything which requires special authorizations only granted to IBM systems programmers. Most of the aftermarket tools not handled by Systems Programming are taken care of by one of the subgroups within Information Resource Management.

This group also works with Database Administration and Applications Development on performance tuning. The hardware/MVS-level monitoring of things like channel traffic and relative CPU usage is its responsibility. Note that the Database Administration group is involved with monitoring the internal workings of DB2 itself. There is some crossover here to the mutual benefit of both groups.

For the DB2 environment, the measure of success is based on various standards.

In addition, Systems Programming is responsible for many aspects of recovery, including full-pack backups and shop-wide disaster recovery.

Computer Operations

Computer Operations starts and stops DB2. It also backs up the system. Note that Operations does *not* do complex recoveries. Either Systems Programming, Applications Development or Database Administration will do those, depending on the nature of the problem.

In addition, Operations migrates application code (COBOL, CSP and so on) from test to production. It also runs production batch jobs and controls all nighttime batch job scheduling. Finally, it controls distribution of paper output.

Information Resource Management

Security Administration

The Security Administration controls RACF and uses RACF to control sign-ons and file-level access. It is the first contact for issuance of all DB2 Grants and Revokes. If no Security Administration staff is available within the required timeframe, then Database Administration is the second contact. To minimize the laborious Grant/Revoke procedure, Glaxo uses RC/Secure from Platinum Technology, Inc. (Lombard, IL) to help automate the process.

Data Administration

Data Administration is responsible for logical data modeling. Glaxo believes you cannot administer data you do not understand and little development work is done without an associated data model passing through Data Administration. This is true not only for data destined for DB2, but also for the VAX Oracle applications and some of the PC development work. As in anything, there are exceptions: for example, ''personal'' databases are allowed but are never modeled by Data Administration.

Also, this group is responsible for data quality assurance. Glaxo's assumption here is that all data is *not* created equal. *Huge* volumes of data are processed every day and you want to keep the size of the Data Administration staff within reasonable bounds. Therefore, a two-phase approach to data quality assurance was developed.

Data Value Analysis

Through a locally-developed algorithm, determine which data is most urgently in need of attention and pursue only the highest-priority subset that can be handled well with existing resources. Data of most value to the company always receives the most effort.

Automation

When a problem area is identified, use the tools at your disposal to minimize the need for manual intervention from the DA group to ensure data quality. An example of a success in this area is the developed referential integrity checker that can be run by the applications developers without intervention from the DA staff. It automatically logs results to a table that can be asynchronously monitored by the DAs. This product is external-table-driven under DB2 Version 1.3 and will be modified to use the DB2 catalog in Version 2.1.

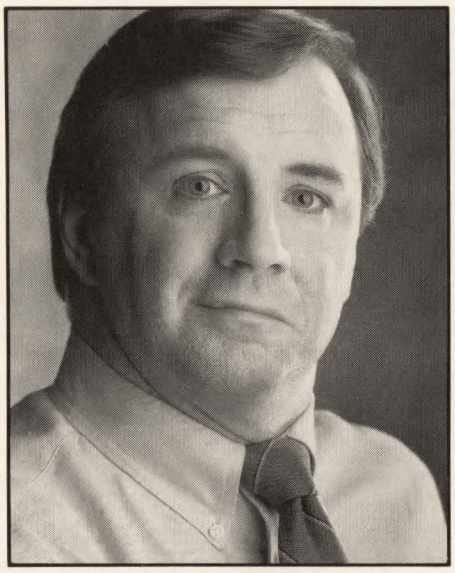

You know your CICS from your IMS...

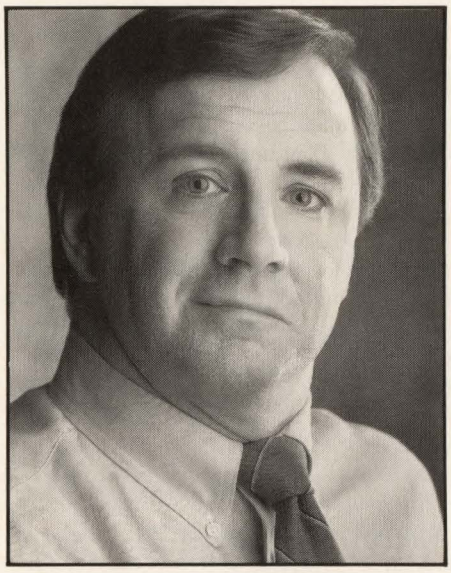

and your VTAM from your DASD.

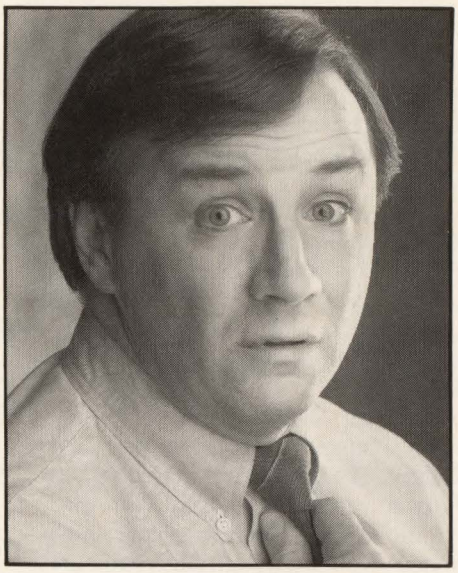

But how about your CAR-RT SORT from your ZIP+4?

We make it EZ as ABC. When it comes to running a DP center, it's a jumble out there. The last thing you need to worry about is managing huge mailings for your company—*and* doing them inexpensively.

Fortunately, there is a solution. Group 1 Software® gives you the power and flexibility to address all mail-related demands through comprehensive software and services—from list management and postal discounts to demographics and personalization. We even offer a variety of consulting services from our expert staff of mailing professionals who will assist you in your data center.

Group 1 is the industry leader in postal and list management software, offering *total solutions* for high-volume mailing needs, plus an unmatched command of ever-changing postal regulations in the U.S. and Canada.

We're so sure Group 1 Software will save you time and money, we'll perform a **free** benchmark test using *your* mailing list. Just return the attached coupon for more information. We'll also send you a **free** 42-page publication on business mailing preparation.

Stop losing Z's over high-volume mailings. Call or write Group 1 Software today.

1-800-368-5806

Yes! I want to take charge of my mailings—PDQ!
Please send me more information about Group 1 Software.
6404 Ivy Lane, Suite 500, Greenbelt, MD 20770-1400

Name _____ Title _____

Company _____

Address _____

City _____ State _____ ZIP _____

( ) _____
Phone

GROUP 1®
GROUP 1 SOFTWARE

MJ89 The Direct Marketing Software Company

CIRCLE #161 on Reader Service Card ▲

Glaxo will continue to use this homegrown tool under the new release because the output from the Version 2.1 CHECK utility is considered to be a little unfriendly for a programmer and because Glaxo is pleased with the logging features built into the RI checker. This product and others which access DB2 directly from the ISPF environment were developed using an attachment facility called RLX from Relational Architects, Inc. (New York City).

In addition, Data Administration controls and supports the corporate data dictionary. The dictionary is used as a repository for entity and attribute information from the data model. It contains cross-references from entity/attribute structures to table/column structures but contains no information about entity-to-entity relationships.

Glaxo also uses the dictionary as a means for creating DB2 syntax. It has been locally customized to automatically construct standard alias names for SQL, CSP and Focus and to generate DB2 DDL syntax, thus allowing easy movement from logical entities and attributes to physical tables and rows.

However, the fact that it is an inactive dictionary makes it less than reliable for monitoring and controlling the DB2 environment once entities are created. Glaxo relies on the DB2 catalog once an entity becomes real, using Platinum Technology's RC/Query to compose catalog queries without bogging the programmers down in catalog structural detail.

Database Administration

Database Administration controls DB2 backup and recovery. Computer Operations in general runs the backups, but the DBA group develops the procedures and has final responsibility for DB2 database recovery in all but full disaster-recovery situations which are handled by Technical Support.

Also, this group guarantees database structural integrity. To ensure the database has no broken pages, invalid index pointers and so on, Glaxo uses Platinum's Database Analyzer to regularly scan database files for problems. This scan is consolidated with the database backup procedures so that only one pass is made through the data for both purposes.

In addition, Database Administration monitors space availability and utilization. Glaxo manually maintains its own DB2 tables in which to record predicted space requirements. Space requirements for all applications are calculated and posted to these tables at the beginning of the development cycle. Then, these predictions are routinely compared against actual usage statistics from the DB2 catalog and historical growth patterns tracked through Database/Analyzer. This approach provides two benefits. The first is the ability to refine space requirement prediction algorithms based on real-world feedback. The second is to help avoid unpleasant midnight surprises and provide Systems Services with adequate DASD requirement information for each budget cycle.

Furthermore, Database Administration is primarily responsible for DB2 performance tuning. (More will be said on this subject later in the article.) Also, it does *some* of the software installation and support. As mentioned earlier, Systems Programming does the bulk of systems-level software installation and maintenance.

Database Administration is responsible for migration of DB2 objects (except Plans) from Test to Production.

In addition, this group handles DB2 support from a programmer or end-user's perspective. It teaches in-house classes how to use DB2 and the related toolset and provides on-line documentation through a locally-customized version of DB2 Guide/Online (Platinum Technology).

It also works with a task force consisting of a representative from each of five major DB2 application areas. The members of the task force are individuals who have shown the ability and desire to provide first-line support for DB2 to members of their areas. The task force members receive training in the use of DB2 and its related toolset beyond that given to other applications developers. Also, they are given low-level access to and are trained in the use of two tools: Omegamon/DB2 from Candle Corp. (Los Angeles, CA) and DB/Optimize from Systems Center, Inc. (Reston, VA), not available to the general programming or end-user communities. The intent here is to broaden the base of DB2 support while keeping a lean and mean DBA staff.

Furthermore, it provides an internal hot line and bulletin board to provide quick support for hot problems. Additionally, all DBAs have an open-door policy.

Database Administration is responsible for migration of DB2 objects (except Plans) from Test to Production. This process is expedited through the use of RC/Migrate from the Platinum Catalog Facility. In addition, it provides a fallback to Security Administration for issuing Grants and Revokes. As noted above, Security Administration has primary responsibility in this area.

One important point needs to be made involving Data Administration and Database Administration. These positions are by nature symbiotic and both positions are filled with people whose goals and experience base extend into the other area. The areas of overlap include implementation of a logical model as a physical database design. Specifically, the DBA needs to have confidence that any logical model produced by the DA will be a reasonably implementable design and the DA needs to feel that the physical implementation will be a fairly pure (if perhaps a *little* denormalized or otherwise altered) rendition of the logical model.

The issue of mapping an altered physical design back to the original logical design has been only partially addressed. Glaxo has designed a database structure to hold both the logical and physical core entities along with some physical relationship rules and to group them as required. But a need that has not been satisfied to date is to be able to explode that database into a diagram.

Glaxo has found no tool to automate that process although a couple are currently being evaluated. For now, the diagrams are drawn manually using a PC-based drawing tool. Glaxo believes a tool may exist that would automatically draw a physical model from the database. Glaxo is less confident there is a tool to do logical/physical maps.

Another overlapping area is in the use of the corporate data dictionary. The DA group is responsible for populating (note: applications developers also assist in the population process) and maintaining the dictionary and the DBA group is respon-

sible for using the entities in the dictionary to generate base DB2 syntax. Once the syntax is generated, the DBA staff will often make changes for the sake of performance, thus subtly altering the physical design to be a little different from the logical model as mentioned above. But since the dictionary is passive, it has no knowledge of the changes unless someone manually retrofits them, which does not always happen. Because of this problem, reports are run regularly indicating differences. These are reviewed by the DA staff and any differences deemed critical are fixed by whomever may be most able to find time to do so.

Furthermore, the Information Asset Directory (IAD) overlaps. This is a tool Glaxo is developing. Only pieces are available now, but the intent is to logically extend the catalog with a DB2 data structure and associated applications replicating the functions needed from a data dictionary and being automatically resynced with the catalog on a nightly basis. Because the IAD is a DB2 database, any information in the catalog will not be carried redundantly in the IAD. The IAD will include the ability to generate language-dependent aliases, map logical entities to physical tables, generate skeleton DB2 syntax and (if we get lucky) someday feed into a tool to draw diagrams.

Application Development

Application development is divided into two types: *traditional* applications and Decision Support Systems (DSS).

The traditional applications are developed using a standard formal project development life cycle. They include both batch and on-line components. The batch parts are usually developed using either COBOL, QMF (IBM), SAS from SAS Institute Inc. (Cary, NC) or Focus from Information Builders (New York City). The on-line components are developed using CSP where possible and CICS COBOL where required.

CSP development is usually done under TSO for production execution under CICS, but there are occasional exceptions in which applications are developed and executed under CICS or developed under TSO for production TSO implementation. To expedite the CSP development process in the TSO environment, Glaxo has written CSP transactions to provide access from CSP back to ISPF. This speeds up the development cycle quite a bit, as getting out of CSP to ISPF, then back in again, is frustratingly time-consuming.

DSS applications are developed using a truncated project development life cycle and only 4GL tools; Glaxo uses Focus, Express from Information Resources, Inc. (Waltham, MA), SAS and QMF. The truncated project life cycle is acceptable because no DSS application actually updates base production data. Data is summarized into intermediate tables in some cases. In other cases, detailed data is retrieved directly. In either case, the DSS application is for *retrieval only.*

A decision-support development team is in place. This group uses Express as its main application development tool. Either data is extracted directly from DB2 tables into the Express database or QMF is used to summarize data from DB2 tables into transient tables in a DSS database from which the data is extracted into the Express database. The decision about summary versus detail is made based on volumes, run-time requirements, number of times summarized data can be reused and

the level of detail required.

One point of interest about the Express product: it consists of both mainframe and PC components with the PC applications using the mainframe DBMS as a data server. This combines the number-crunching capability of the mainframe with a user-friendly interface on the PC.

Some DSS work is done directly with QMF and more will probably be done in that environment as the product matures. As with the Express applications, sometimes input comes directly from the detail tables and sometimes summary tables must be created first. Also, some end-user types use Focus or SAS, with the same basic summary/detail data-access structure as was mentioned for Express and QMF.

Backup and Recovery

DB2 has a sophisticated tablespace backup/recovery mechanism. Tablespace-level backups (called image copies) can be taken at any time, even with DB2 running. The image copies can be of the entire tablespace (full) or just of pages changed since the last image copy (incremental). Backup status information is stored in a DB2 catalog table named SYSCOPY. DB2 also keeps a recovery log of all database changes and records information about the log in a directory table named SYSLGRNG and in the BSDS Bootstrap Dataset (BSDS).

The standard DB2 recovery mechanism is highly automated and can intelligently use full-image copies, incremental-image copies and log files to recover a tablespace back to either a given image copy or to the last COMMIT point. With some assistance from the DBA staff, recovery can optionally be done to a given point in the log; but doing so is non-trivial because the process of selecting the appropriate RBA is difficult.

Note that this entire process operates at the DB2 *tablespace* level and the process runs too slowly to allow all tablespaces to be backed up together in the available batch window. IBM provides no high-speed, full-pack backup facility that automatically updates the SYSCOPY table with information required to allow the full-pack backup files to be automatically included in the DB2 recovery process on a file-by-file basis. Additionally, IBM provides no way to manually flush the logs to tape and update the appropriate system tables to reflect current status.

In order to operate within the constraints just described, there are two levels of backup. One is a weekly full-pack backup of all DB2 packs using DFDSS (IBM). This happens every Sunday night with DB2 down and is intended to be used for *last-resort disaster recovery only.* These are *not* image-copy backups, and the system control tables described above are not updated. While these full-pack backups provide total recovery capability on Monday morning (assuming nothing has been run after the backups), their value diminishes rapidly throughout the week.

Since DB2 roll-forward recovery is catalog/directory driven, this full-pack backup alone cannot be used as a basis for point-in-time recovery. The log tapes and image copies which are created as DB2 runs throughout the week are in no way synchronized with Sunday's full-pack backup. If the full-pack backup is restored, then the catalog tables controlling recovery are also restored, functionally disabling (as far as automatic recovery is concerned) all image copies and log tapes created after the full-pack backup was taken.

In developing second-level backup procedures, Glaxo had some additional self-imposed constraints.

Since the full-pack backups really are not worth much in terms of DB2 recovery, the second level of backup has to serve two purposes. One purpose is to provide disaster recovery capability up to the end-of-work on the day before the disaster. The second one is to provide the ability to recover a given tablespace up to the last COMMIT point in a non-disaster-recovery situation.

In developing the second-level backup procedures, Glaxo had some additional self-imposed constraints. One was that database backups should happen under Computer Operations control without the DBA staff having to be involved in day-to-day routine.

Secondly, the standard DB2 toolset should be used as much as possible. Since the backup/recovery process is so integrated with DB2, Glaxo felt the only way to be insulated from future software changes was to use IBM's tools and to play by the rules.

A third constraint was to not back up anything unnecessarily. Glaxo does a lot of application-level image copies as part of normal nighttime batch processing to allow quick recovery should the related batch job fail. This backup is often considered to be sufficient, subject to evaluation by the rules that follow.

The last constraint was to minimize disruption to applications from the backup process.

The procedures Glaxo developed depend on the DB2 intelligent recovery mechanism and the availability of all log tapes and image copies required for recovery.

To ensure that there are image copies to recover from while abiding by the constraints above, the following automated backup methodology was developed.

To begin with, back up tablespaces by database level. This requires that databases are organized with backup in mind. The decision about which tablespaces go in which databases is based on relationships from the logical model. Then those are subdivided based on group stop/start requirements and (if necessary) subdivided again for backup. Note that the system work database DSNDB07 and any personal databases are not included as part of the automated process.

Since there are always exceptions to everything, create a table containing the names of tablespaces to be exempted from the automated process and never automatically back up those tablespaces. This was required for tables which are loaded from other data sources and never changed and large tables whose automatic backup would take so long as to disrupt other applications. Those tablespaces need to have manual procedures developed to handle them, generally involving incremental backups with regular consolidation points.

It was also required for permanent *temporary* tables retained as transient workspace and anything else that might be chosen to not back up automatically at a given time for whatever reason.

If a tablespace has never been backed up (no entry in SYSCOPY table), take a full-image copy of it. If a tablespace has not been backed up in the last 15 days, take a full-image copy of it. Glaxo does this by executing a small program at the beginning of the backup job that calcu-

lates the appropriate date and posts it to a one-row table that is then used to compare against the dates in the SYSCOPY table. DB2 date arithmetic could not be directly utilized because the date column in SYSCOPY is in YYMMDD character format.

If more than 25 percent of the pages in a tablespace have changed since the last full-image copy, take a full-image copy. If batch run-time requirements prove to be prohibitive in the future, then Glaxo may use incremental copies followed by later consolidation, but there is no need to do so at this point.

Finally, if none of the criteria above is met, do not take a backup. Since this job is run after other nightly processing, any table that was backed up as part of an earlier batch stream will not be redundantly backed up again unless more than 25 percent of the pages in its tablespace have changed since the earlier backup.

The methodology has been implemented through the use of Database Analyzer, a tool intended for analysis of DB2 data structures. It allows specification of databases to be analyzed, then analyzes the individual tablespaces within the database and applies user-customizable criteria against the results of that analysis. Without this tool it would not be possible to implement the ''25 percent of pages changed'' criteria specified above. All of our criteria are specified in a customized query, the results of which cause Database Analyzer to submit a batch job to back up selected tablespaces.

A positive by-product of this process is that the results of the analysis performed by Database Analyzer are written to DB2 tables from which further information can be extracted as desired.

To ensure availability in the event of a disaster, all nighttime DBA-generated image copies and copies of each log tape (Glaxo does dual logging) are sent off-site.

There are a couple of miscellaneous points of interest. One is that the analysis/backup jobs are run on a recurring seven-day cycle. Job submission is through an automated scheduling package, CA-Scheduler from Computer Associates International, Inc. (Garden City, NY).

In addition, Glaxo uses the MODIFY utility to ensure that the SYSCOPY table has only a rolling 45 days of information for all tablespaces. Therefore, long-term archival image copies must be manually recovered using the DSN1COPY utility that works without requiring an entry in SYSCOPY.

Performance Analysis

Glaxo's three performance analysis tools are a real-time control block monitor, an SMF reporter and a database structure analyzer.

A real-time control block monitor, Omegamon/DB2, allows spot-analyzing of immediate performance problems. An SMF reporter, DB2PM (IBM), derives information from SMF statistics and accounting information that is accumulated from DB2 at regular intervals. This tool helps analyze long-term trends and more effectively plan future resource requirements.

The last tool is a database structure analyzer, Database Analyzer, that is executed as part of the routine backup process. The output from the backup analysis phase is recorded in time-stamped rows in DB2 tables. From there the data can be used for both immediate problem diagnosis and historical trend analysis. This information is used to routinely monitor fragmentation, poor index distributions, file growth patterns and so on without having to redundantly run separate analysis jobs.

The DBA group analyzes information gained from these three tools and then works with Systems Programming and/or Applications Development to affect changes as required. The intent of this cross-group effort is to tune MVS (CICS and so on), DB2 and the application programs in such a way as to provide maximum shop throughput while ensuring that computing resources are dedicated to DB2 applications in a manner consistent with the value of the information returned by those applications.

To evaluate success in performance tuning, some metric must be established. Glaxo's measurement criteria differ, depending on the type of activity being addressed.

For *transaction processing,* the major metric (like everyone else's) is response time on a loaded system. Instantaneous response is preferred but (to be tactful) it is not always achieved. However, instantaneous response is not the whole story. *Acceptable* response times are different for different transactions, varying based on the *perceived value of information received.*

Glaxo aims for instantaneous response and tries to never exceed three seconds. However, living in the real world with real physical constraints, not all of the transactions are instantaneous. Still, there has been success in returning *acceptable* response times. Too little information is not of much value, regardless of speed.

Realizing that relational technology offers some opportunities to break the constraints of traditional application design techniques, Glaxo is experimenting with data driven application development techniques. The IRM group has designed an MIS Administrative Database that includes all of the entities used to administer menus, program linkages, people, userids, application level security and relationships between entities. A better balance between speed of development and transaction execution versus information presented is the goal.

One way data driven design techniques were used to shorten the response times perceived by the user was to reduce the number of screens a user needs to respond to by eliminating the need for functionally oriented menus. The MIS Administrative Database applications are structured to reflect the logical data model. Two types of screens are used: an entity maintenance screen and an entity list screen. Whenever a foreign key appears in an entity maintenance screen, it is identified with a > character. The user can overtype the > with a command character: (I)nquire, (A)dd, (C)hange or (D)elete and be transferred to the entity maintenance screen identified by the foreign key value. There is no need to return to a menu and write down a key to move to another entity maintenance screen.

Additionally, if the user types a value in a foreign key field, including mask characters (_ and %), the user will be transferred to the entity list screen where (s)he can select the row containing the key desired. Upon return to the original transaction, the selected value is inserted in the field from which the query originated.

Each entity maintenance application screen displays a list of relationships in which that entity participates. The user can select one of the relationships and *explode* into the related entity list screen to see only the rows related to the *parent* or *owner* entity. These entity list transactions can be entered from a menu or any number of *owner* entities, so they are reusable and are able to serve many purposes. With these techniques, a user can enter the database and navigate through it using a logical data model diagram as a road map.

The relationship navigation is currently hard-coded in the applications. Glaxo is

See Glaxo *page 91*

The Importance Of Service Level Objectives

Setting SLOs removes the moving target that usually plagues performance and capacity analysts.

By John E. Fair

On an increasing basis, large computing centers are realizing the importance of establishing internal Service Level Objectives (SLOs). Management of these companies is realizing that without these formalized objectives, it is practically impossible to measure how well they are really doing. The SLO is no longer the buzzword of the eighties but a necessity.

Also, a growing number of companies are not only establishing internal SLOs, but also they are drafting formalized Service Level Agreements (SLAs) with their users. These agreements typically guarantee that data processing departments will provide a specified and stable level of service for high percentages of the user's workload.

The intent of this article is to discuss key issues that should be addressed with respect to setting and managing SLOs and attendant SLAs in an IMS environment.

The SLO is an internal commitment to process a high percentage of a user's transactions within a predefined interval of time by using a predetermined amount of computing resources. As simple as that sounds, the SLO should be the basis for operating any interactive system environment. The SLA or the formalized commitment to the SLO between Information System (IS) management and user departments should be considered following general IS acceptance of the SLO.

The SLO provides a fixed target for capacity planners and performance tuners. Without objectives, any performance tuning or capacity planning effort is difficult because of constantly moving targets. As a general rule, it is also not satisfactory to set a global SLO such as, "We will process 95 percent of all IMS transactions within five seconds." Reasons for not having this type of SLO will be evaluated later in this article.

Prerequisites For Establishing SLOs

Before IS can establish SLOs, it must establish its baseline and understand the workloads.

- As part of any evaluation of SLO concepts, IS must first *establish its baseline*. The organization must determine whether there is adequate capacity to allow it to guarantee better service than is currently being provided.

Establishing the baseline involves a number of issues. It may also involve a number of different groups within the IS organization. Unless the organization is planning to promise status quo service, it will need to fully understand its current performance and capacity profiles. Can I provide better service? This question cannot be answered without substantiating data.

What parts of the organization are affected becomes the issue. The answer is simple. Any part of the organization that contributes to meeting or to not meeting the SLO. Simply stated, if an SLO includes DASD service, the DASD profile must be clearly understood.

It is up to IS management to *validate profiles* of *workloads* for which SLOs are being set. Each of the following items needs to be considered on its own merit.

Characterize Workloads

There are many articles written on workload characterization and qualification. I will not presume to describe the mechanics of workload characterization but merely point out that it should be done.

Define As Few Discrete Workloads As Possible

Tracking transaction clusters is more appropriate than transaction codes with a cluster being defined as a discrete grouping of transactions that share a near-common profile. Clustered transactions should be approximately the same size and do approximately the same amount of work.

VPS

Tames The Jungle of MVS Network Printer Support

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).

1700 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 modifications. NO JES maintenance.
- ✔ Automatic forms control, full FCB support, dial-up PC printers, printer pooling and "hot" printers are all supported.
- ✔ Full screen "ISPF-like" command interfaces 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

ATTENTION: Marketing Department

LRS *Levi, Ray & Shoup, Inc.*

Specializing in Computer Systems Software

2387 West Monroe
Springfield, Illinois 62704
(217) 793-3800 ◆ Fax (217) 787-3286

® SASWTR is a registered trademark of SAS Institute, Inc., Cary, N.C.
® ROSCOE is a registered trademark of Applied Data Research, Inc.

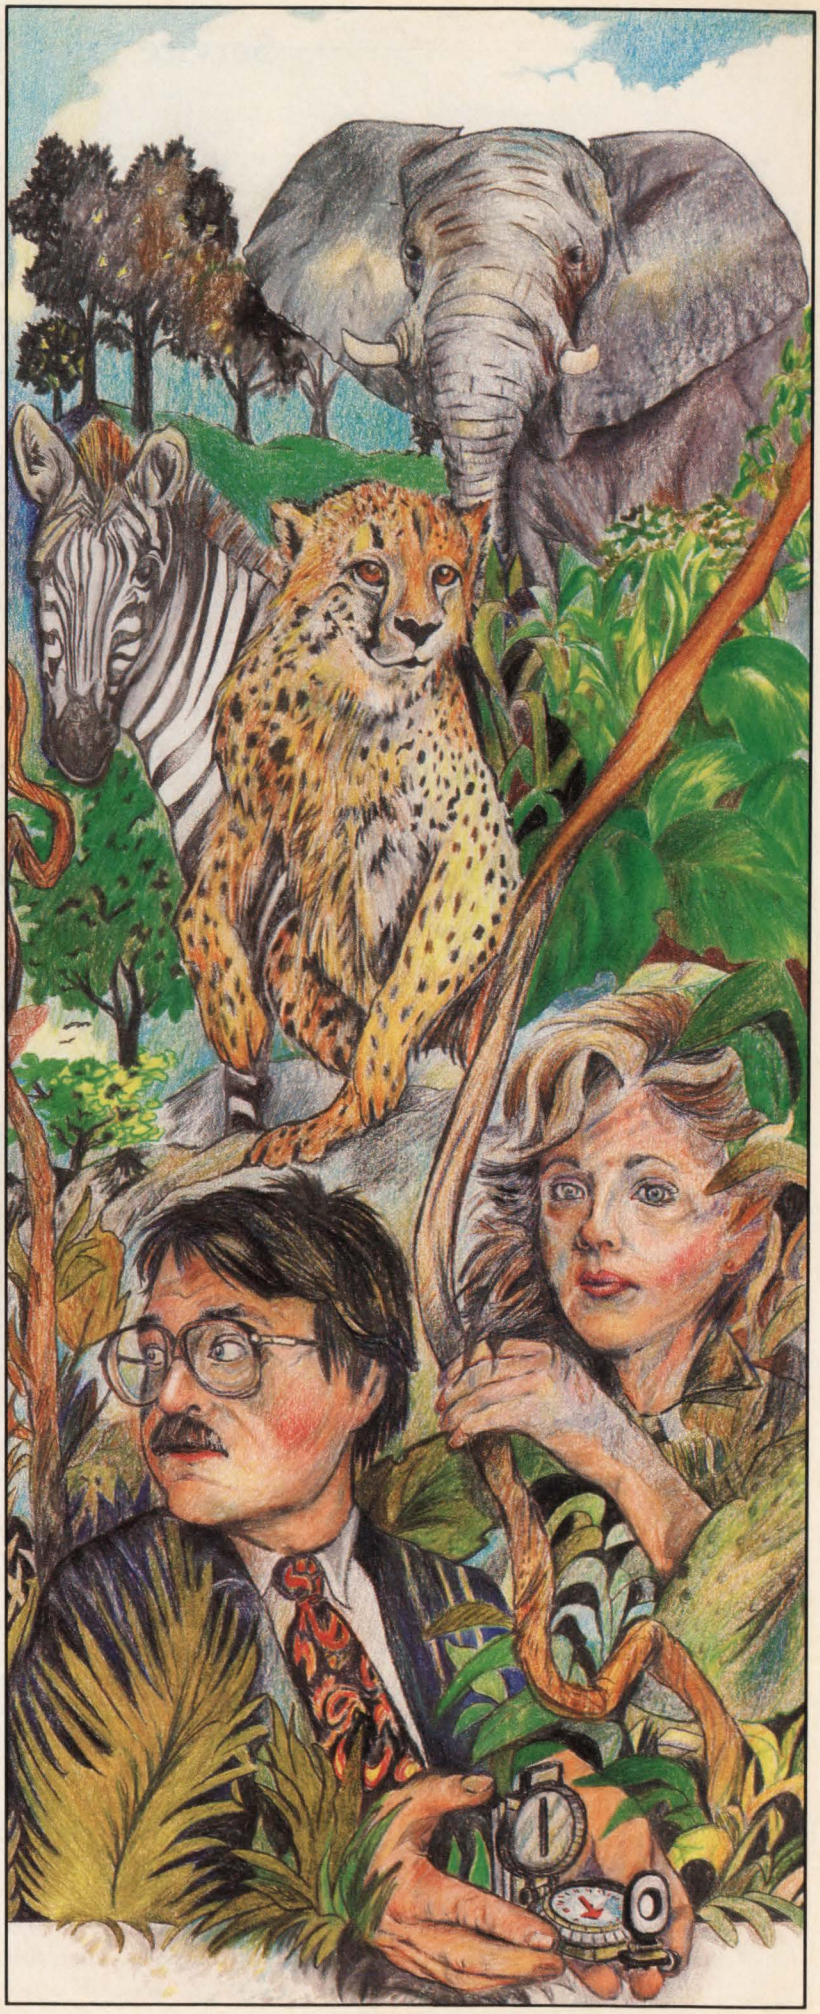

Transaction clusters will be discussed in more detail in the next section.

Understand The Variability Of The Workloads

SLOs usually apply to all occurrences of the same workload and assume that there is little variability across different executions of the transaction or job. For SLOs to be effective, different executions of the workloads should be repeatable with little variance from one execution to another. If there is much variance, the SLO should be built with the worst case scenario in mind. This is especially true in cases in which SLAs will be negotiated based upon the SLOs.

Establishing A Baseline

Key to the successful implementation of SLOs is the accurate establishment of a baseline. The following procedures should be followed.

Develop Clusters

A typical IMS shop has hundreds of transaction codes defined. Because of the mere volume of defined transactions, it is usually not practical to manage SLOs at the transaction level. Unless you represent a small IMS shop, do not try to set SLOs for each discrete transaction code. You will be better served by reducing the number of defined transaction codes to a manageable number of clusters.

Clusters typically encompass a group of transactions that share a common profile or represent transactions from a specific operational unit within a company. I prefer the former approach because there is usually less variability between transactions within the cluster with the latter.

The number of defined clusters will vary from company to company, but I believe that less than 20 is a manageable number for most companies. Having a large number of clusters becomes unwieldy and usually indicates that a less than adequate job has been done during transaction profile analysis.

Model The System

Modeling the system is one of the best ways to define clusters. Depending on the modeling tool chosen and whether it is a simulation or analytic model, you must do slightly different things; but several rules apply to any modeling effort. A partial set of these rules is covered below.

- Choose your sample(s) carefully. It is important that you sample from peak but non-thrashing periods of time.
- Be sure the workloads are representative. It does little good to model the system if the samples do not reflect the real world. I suggest that you use all available tools to ascertain from which points in time the best combination of transaction and batch workload may be available. Though the cadre of tools will vary from company to company, I believe that a combination of tools including RMF from IBM, IMF from Boole & Babbage Corp. (Sunnyvale, CA), Omegamon from Candle Corp. (Los Angeles, CA), the MICS IMS Analysis Product from Legent Corp. (Vienna, VA) and SAS from the SAS Institute (Cary, NC) will be useful.

To establish a baseline, IS must decide if there is adequate capacity allowing a guarantee of better service.

- Be sure the workloads to be modeled are as complete as possible. This will almost certainly mean taking more than one sample; therefore, you will run and calibrate the model more than once.

Modeling the system accomplishes several basic objectives. It allows the person doing the modeling to calibrate a model to the *real world*. This will be valuable when playing *what if* games and when projecting the effects of adding new workloads to the current system(s). It will provide information from which the clusters that have been discussed can be defined. In order to build these clusters, look at as many of the discrete workloads as possible at least once.

Cover Your Backside

There will be outliers that do not quite fit the mold but that are not far enough out to be placed into another cluster. There will be times when the system does not run well for unforeseen reasons. The number of these conditions must be kept at a minimum but they will exist.

Because of these conditions, SLOs should *never* be based upon guarantees of 100 percent of the workload falling within the objective. The percentage of variability will depend on how good the clusters are, how repeatable the workloads are and on the availability of the system.

The SLO should never be for less than 90 to 95 percent of the time.

Tracking And Monitoring SLOs

The key to the success of any SLO program must include the following key points.

Monitor Successes And Failures

Reports should identify the clusters and report how well objectives are being met. Using exception reporting seems to be the best approach to management of SLOs. Exception reports should identify the outliers. They should also report any upward trends in resource utilization that may in the future affect missing the SLO.

Track Changes Made To SLO Applications

Anytime a change is made to an application that is subject to the SLO, the effect of the change should be measured. This is just as true for the one line change as for the rewrite of the program. Since a single application can adversely affect the performance of an entire IMS system, the old adage, "But I only changed one line of code," cannot allow changes to go invalidated.

Classify And Validate New Applications

Once an enterprise adopts an SLO policy, all new applications placed on-line should be subject to the policy. The classification can be accomplished as part of a multi-step process which includes five steps.

- Review the design of the new applications to determine that the overall design conforms to standards for SLO applications.
- Test the applications thoroughly, running monitors and diagnostic aids extensively. During this testing phase, accurate transaction profiles can be developed.
- Project transaction volumes. In order to determine the effect of a new application, its transaction volume must be known. A great deal of time should be spent on analysis of this issue, since

GIANT

IS NOW ON LINE

The Global Information Access Network

Making decisions about hardware, software and data processing services is a difficult task. Just trying to identify the available products can be an enormous dilemma.

The time has come to benefit from the same technology you work with day-to-day. Explore the numerous products and services available from the comfort of your office. GIANT provides a complete data base at your fingertips through our dial-up network. And best of all, it's FREE!

Free Access — No User-id Required — No Password — On-line Instructions.

Dial-up access is available from your PC. Use any of the popular communications packages that emulate a VT100 terminal (or use a VT100 terminal). We have successfully tested using PROCOMM Plus™, CROSSTALK™ and KERMIT.

2400 → (414) 546-8480
1200 → (414) 546-8492
Terminal type VT100
Line speed . .2400 or 1200
Parity. . . . . .None
Data bits 8
Stop bits 2

Dial-up is available Monday through Saturday, 24 hours a day. (BISYNC and SNA dial-up support coming soon)

If a dial-up terminal is not available, then FAX your request to us. We will search our data base and FAX a list of products or services to you.

(414) 546-8499 (FAX)

FAX is available 7 days a week, 24 hours a day.

Or you may call us direct, and we will search our data base immediately to find the products and services you need.

(414) 546-8496
(Voice Lines for Personal Service)

Personal assistance is available Monday through Thursday from 9:00 AM to 1:00 PM Central Time.

Software Manufacturers!! Hardware Manufacturers!! Service Suppliers!!

Advertise your products and services on GIANT today. Put your message into every data center across the country . . . everyday. Join with GIANT and become part of the GIANT team: we're making the marketing media of the future a reality today!

Give us a call and we will send our GIANT information kit to you immediately.

Productivity Systems, Inc.
P.O. Box 20958
Milwaukee, WI 53220
414/321-8688

PROCOMM Plus is a registered trademark of DataStorm Technologies, Inc. CROSSTALK is a registered trademark of Data Communications Associates, Inc.

CIRCLE #135 on Reader Service Card ▲

missing estimates by only a few percentage points can later have negative impact on the system.

- Find an existing cluster that fits the profile of the current workload. I prefer to add a cluster only when absolutely necessary, since adding a cluster can distort some models.
- Add the projected workloads to the cluster(s) and run the model again. Since the model has already been calibrated, you will really be playing a *what if* game at this point. The output from the model will allow you to project the effect of the new application with a great deal of accuracy. It will also allow you to predict how responsive the new application will be.

Develop Standards And Insist On Strict Conformance

While this is the most difficult of the issues at hand, it is also one of the most important. Conformance to strict standards set for each of the clusters will do much to ensure standardization of workloads. Enforcement of these standards will go a long way toward ensuring more repeatable transactions, more system throughput and a better running IMS system.

Basis For Setting SLOs

SLOs can be set on a number of variables. In the following section several are discussed.

CPU Time

CPU time is without question the most representative of all measures. Assuming little variability between different executions of the same IMS transaction, the CPU time should be approximately equivalent across executions.

Depending on the collection tool, different measures of CPU time are available. IMF users have excellent measures of CPU resources consumed. In addition to measuring message region TCB time, IMF measures CPU time used to complete tasks that are usually considered to be *system overhead,* things that are generally not measured.

Input/Output Operations

I do not recommend basing IMS SLOs on I/O measures for the following reasons. IMS buffer pools are typically large, shared pools that are used concurrently by a number of applications and managed by IMS. Since writes are typically asynchronous events that occur when space is needed in a buffer pool to satisfy a request for INPUT data, they tend to not be repeatable.

Even though reads are typically more synchronous, they should also not be used as the basis for setting SLOs. I believe that the number of I/O operations required to satisfy a read request is not predictable enough to justify using this measure as a basis for setting SLOs. For instance, a recently reorganized database may require far fewer I/O operations per read request than one that has not been reorganized for several weeks. One execution of a transaction may require I/O and another may not require it because data may have already been in the buffer pool.

Setting SLOs removes the moving target that usually plagues performance and capacity analysts and should be done with or without SLAs.

Database Calls

Another basis for setting SLOs is database call counts. One frequently used objective is *if a transaction uses less than one-half of a second CPU time and 30 or less database calls, it gets one second response time 95 percent of the time.* While this is certainly better than nothing, I suggest that this approach is seriously flawed. A GET UNIQUE call may require far more I/O operations than a GET NEXT. An ISRT call may require more I/Os than a replace.

If database calls are to be used for setting SLOs, I recommend using one of the two following techniques. The first of these includes weighting the calls by type of call. Using this technique, a GET UNIQUE will have a higher rating than a GET NEXT. I have no recommendations here with respect to weights since the variance is highly database and application dependent.

This technique may also be flawed, at least to the extent that all calls of a specific type would weigh the same. This problem can be countered by basing the calculation on *probability of I/O.* Using probability of I/O, the weights assigned to each call type consider access method, number of secondary indexes and other factors that may cause the projected number of I/Os to vary not only across call types but across databases.

For instance, an application using a VSAM KSDS with several secondary indexes would not have the same SLO as an application using HDAM/OSAM with an efficient randomizing module.

Message Length

Another frequently used measure is the length of the input and/or output message. The data sources most often used for this measure are the type 01 and type 03 records from the IMS System Log Dataset (SLDS). These numbers may or may not represent actual message lengths, depending on whether Message Format Services (MFS) is being used. These measures represent characters that are transferred to and from message processing program I/O areas and do not contain 3270 formatting data streams. Whether these measures warrant further consideration is highly dependent upon individual preference.

Summary

A number of issues have been addressed that should be considered prior to establishing a program that includes setting application SLOs.

I strongly believe that IS management in practically all categories of computer processing should insist on SLOs being established regardless of whether there are formal SLAs with their users. Setting SLOs removes the moving target that usually plagues performance and capacity analysts. Instead of waiting for the phone to ring, applications can be measured against specific objectives. Management of computer workloads is put back where it belongs — with IS management.

ABOUT THE AUTHOR

John E. Fair is President of J. Fair Systems, Inc., a computer consulting firm specializing in DB/DC system and performance consulting. He is the author of more than 20 published articles on IMS, DB/DC systems and relational database products. The firm is headquartered in Centreville, VA, (703) 830-2357.

ANNOUNCING

A New More Productive Era in Systems Programming

The Systems Programming Environment Only in the SAS/C® Compiler

Until now, higher-level languages just couldn't hack it in the systems programming world. Too many issues stood in the way—inefficiency, poor access to low-level system services, bulky and intrusive library requirements, and inflexibility in addressing the IBM 370 architecture.

But now you can write systems-level routines faster and maintain them better than you ever could with assembler—with the SAS/C compiler's exclusive Systems Programming Environment (SPE).

SPE is an extension to the C language that greatly simplifies the coding of user exits, tools, and utilities for JES2, VTAM, CICS, TSO, GCS, and other systems software. Included are support routines that allow you to write and execute C programs and a compact runtime library that features both general purpose and system specific functions for memory management, interrupt handling, low-level I/O, and more. There's also a utility that translates assembler DSECTs into C structure definitions—an enormous time-saver when you're writing programs that interface with assembler.

Together these tools provide a freestanding C environment designed to interact with the operating system the same way assembly language programs do. With SPE, your C programs can:

■ call existing assembler routines ■ generate any machine instruction in-line ■ easily access system data and control blocks ■ exploit BSAM or CMS file system I/O ■ process asynchronous events and interrupts ■ directly use SVCs and DIAGNOSEs

Then, at compile time, the SAS/C compiler's global optimizer will compress your code to produce routines that rival assembler for speed and efficiency.

With frequent updates and knowledgeable technical support—both provided free—the SAS/C compiler is the best investment you can make toward greater systems programming productivity.

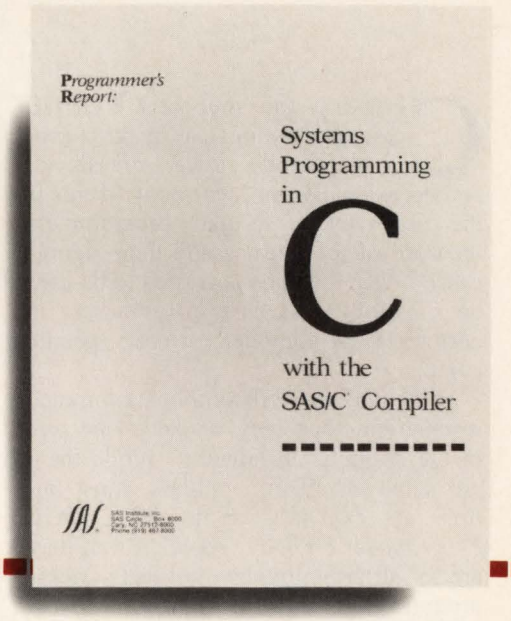

Learn More in a FREE Programmer's Report

To find out more about systems programming with the SAS/C compiler, simply clip the coupon below and mail today. We'll send you our new **Programmer's Report: "Systems Programming in C"**. Or call us today to find out how you can receive the SAS/C compiler for a free, 30-day evaluation.

SAS Institute Inc.
SAS Circle □ Box 8000
Cary, NC 27512-8000
Phone (919) 467-8000
In Canada, call (416) 443-9811

The SAS/C compiler runs under MVS (370, XA, and ESA) and VM/CMS on IBM 370/30XX/43XX/937X, and compatible machines.

___ **Yes, send me a FREE copy of "Systems Programming in C".**

___ **Contact me with details of a FREE 30-day trial of the SAS/C® compiler.**

Mail to: SAS Institute Inc., Attn: ME, SAS Circle, Box 8000, Cary, NC, USA, 27512-8000. Telephone (919) 467-8000.

MFJ889

Please complete or attach your business card.

Name _____

Title _____

Company _____

Address _____

City _____ **State** _____ **ZIP** _____

Telephone _____

Hardware System _____ **Operating System** _____

COBOL COMPILER OPTIONS

Understanding Your Choices

By Harvey Bookman

Choosing the incorrect COBOL compiler options can make a program compile slower, run slower and occasionally run incorrectly. If this is the case, why are so many programmers unaware of the options and their significance? Were the options created to be used on a case-by-case basis or as a way to customize the compiler for each specific company?

If the options were simply set for each installation, then they would all be part of the compiler installation. While there are some installation options, most options are explained in the *VS COBOL Programmer's Guide*, showing that they are for all programmers, not just systems programmers. It is unfortunate that most programmers have accepted compilation procedures as "canned." Whatever their current project uses as a standard compilation procedure is considered "proper" and changes are not made to it. Programmers often feel there is no need to learn the options since a procedure is already set up. I have worked on COBOL projects on which I was literally forced to use standard compilation procedures even after proving their inefficiencies. They can slow down the time of each compile, making a programmer less productive as well as force options that are not optimal for a programmer's debugging ability.

When using the COBOL compiler options, there are some general rules to follow. Some compiler options must be set depending on specific programming products used with COBOL. For instance, CICS requires some options while restricting others, while the IBM interactive debugger requires the compiler to be informed through a compiler option. However, many options depend on subjective factors such as debugging ability (for example, the programmer can understand a Procedure Map) and debugging methodology (for example, the programmer reads dumps rather than recompiling a program with READY TRACE and rerunning). Still others depend on whether the compilation is being done to find compilation errors or for a completely different purpose like creating a listing for documentation.

Choosing an efficient set of options can decrease the compilation time of a program by more than 50 percent (see Figure 1). As will be shown in this article, the 50 percent savings of time in compilation can only be realized at certain times. However, any savings of CPU processing time and elapsed time will save money and increase throughput.

Each compiler option has a default. Some are set to a value if not specified (for example, SIZE = 128K); others are set to either one specific setting or another (for example, QUOTE/APOST, FLAGE/FLAGW). Most are either set on or off depending on whether or not the option is preceded with 'NO' (for example, PMAP/NOPMAP).

Options Guiding The Compiler

One of the least used and most efficient options (from the point of decreasing compilation time) is SYNTAX (the default is NOSYNTAX). It tells the compiler to check for syntactical errors and to not produce object code. This option should always be used when a program is compiled the first time. In the rare case that a program compiles without errors the first time, usage of this option will require the program be compiled again without the option. However, it is rare that a programmer writes any reasonable sized program without at least one compilation error.

The SYNTAX option instructs the compiler not to produce a text (object) deck. There is another option, CSYNTAX, that tells the compiler to proceed

FIGURE 1

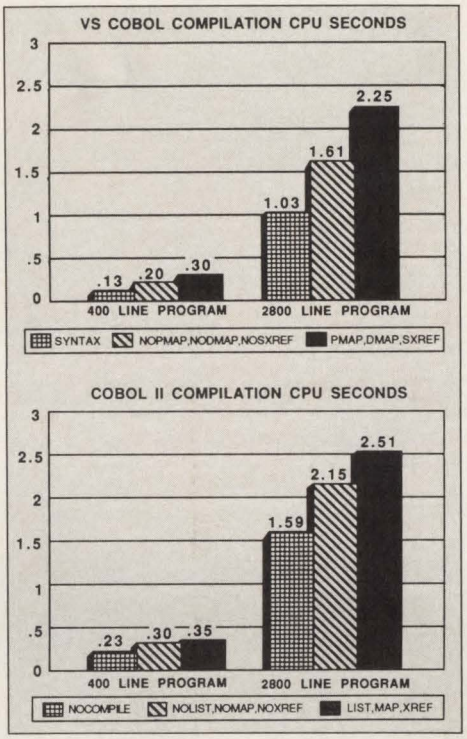

as if it were not doing a SYNTAX compile but as soon as an error exceeding a 'W' or 'C' level is encountered, switch to a SYNTAX compile. While CSYNTAX is not quite as efficient as SYNTAX, it should be used after a large number of compilation errors have been fixed in a program and you are unsure as to whether or not the compile will be error free.

The LOAD and DECK options often cause confusion to programmers. They each instruct the compiler to produce an object deck. LOAD causes the output to go to the dataset specified in the SYSLIN DD statement while DECK causes the object deck to be written to dataset defined in the SYSPUNCH DD statement. LOAD is the default and JCL procedures take the output from the SYSLIN DD statement dataset as input to the linkedit step. NODECK is a default since its output is only used for backup purposes such as when a card deck is punched.

The SIZE option specifies the amount of storage to be used for the compilation. It can be specified in bytes (SIZE=131072) or in kilobytes (128K). A larger SIZE is generally more efficient and will usually result in a faster compilation. This is because the buffer space for the compiler work files is increased and increased storage lowers or eliminates I/O for the dictionary created for the names used in the program. For large and/or complex programs, a SIZE larger than the default (128K) may be necessary. Although it is not clearly stated in the manual, when a SIZE is specified greater than the storage available, the maximum available storage is used. This is why the option is sometimes coded as SIZE=9999999. (In COBOL II the maximum storage size for compilation can be requested with the SIZE(MAX) option.)

The QUOTE option indicates that the double quote ('') is used to delineate literals and constants within the code while APOST informs the compile that the apostrophe (') will serve this function. While APOST was the default prior to VS COBOL, the QUOTE option has been the default since. Usually either the QUOTE or APOST option will be set as a standard for a company and should not be overridden except in special cases. Company-wide copy members and database dictionaries will create data in either one format or the other forcing programmers within a company to be consistent in their usage

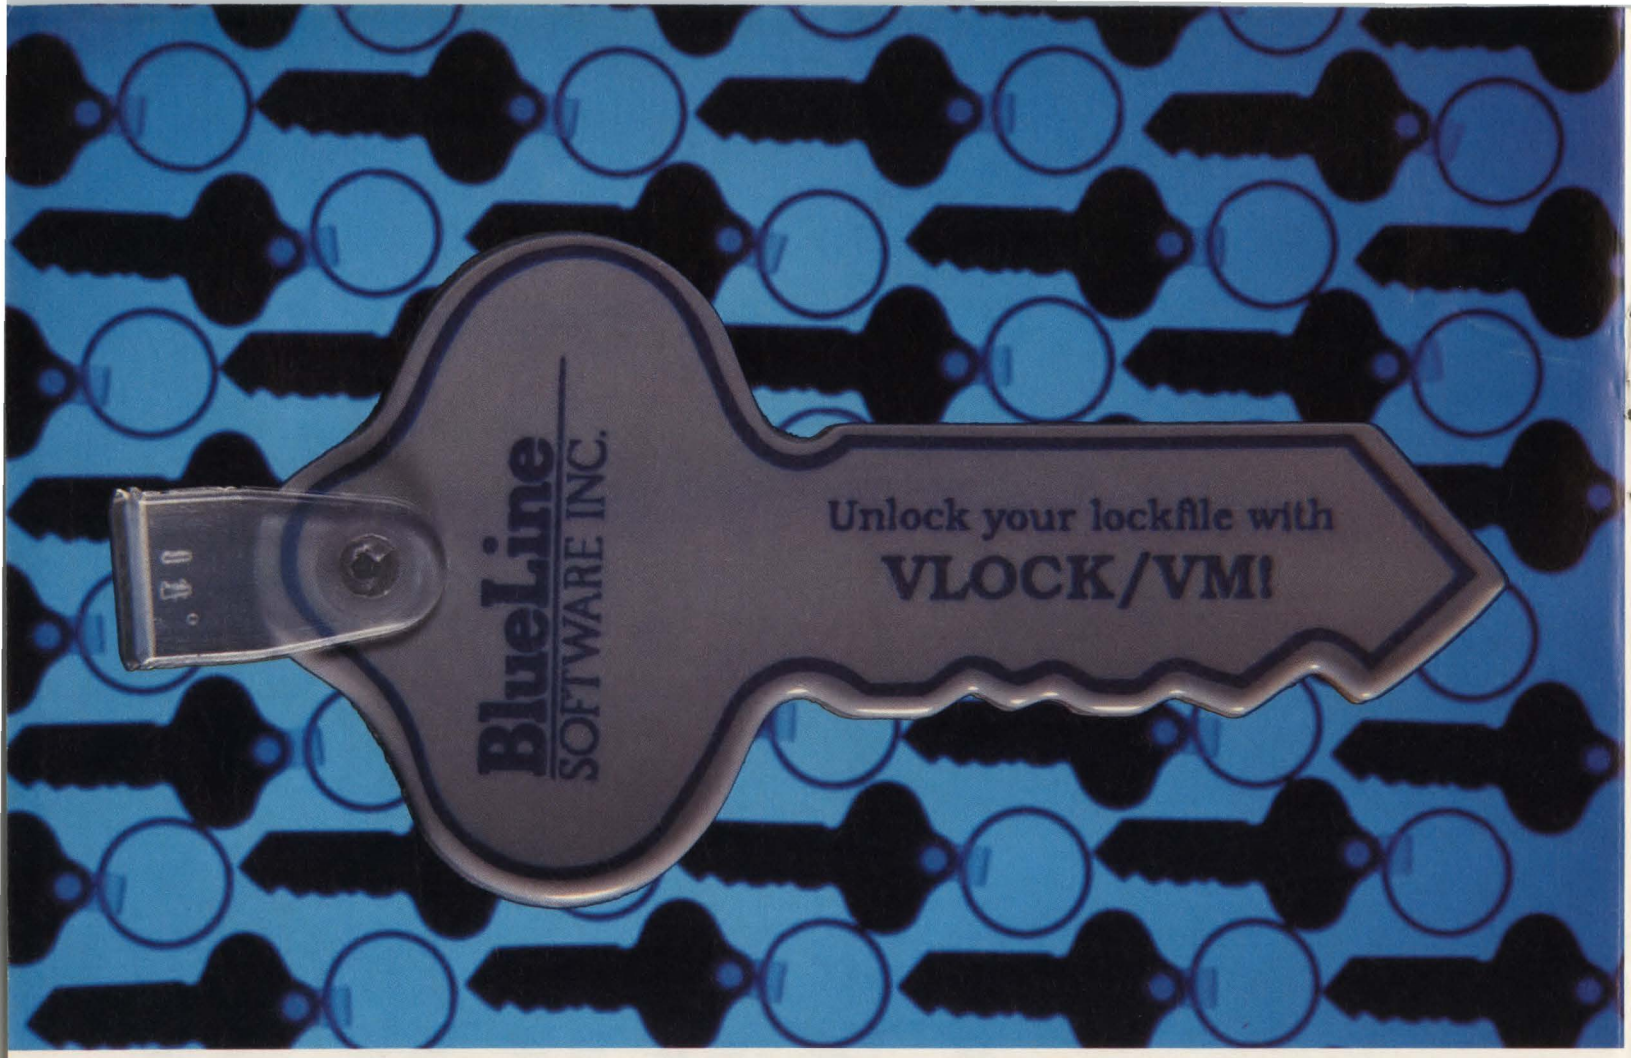

The key to improved throughput.

Do you have multiple VSE guest machines running under VM? Would you like to be able to share DASD files between multiple VSE guests - without the performance degradation of the VSE lockfile?

VLOCK/VM is a VSE lockfile manager that increases the throughput of your system by eliminating I/O requests to the VSE lockfile. All lockfile requests are intercepted and fulfilled in real storage, solving the problems associated with lockfile performance, including device busy conditions and channel contention.

Our customers report that VLOCK/VM has saved them up to 25% of all I/Os performed each day! And since you'll now be able to share DASD files between VSE virtual machines without reducing performance, you may regain some DASD space.

VLOCK/VM is helping more than 200 installations solve their VSE lockfile problems. And VLOCK/VM earned a 1989 ICP Million Dollar Award, further attesting to its leadership position in the marketplace.

So, call us today for a 30-day evaluation at ***800-826-0313 (612-542-1072 in Minnesota)*** *or complete and return the attached postage-paid reply card. You'll be glad you unlocked your lockfile's performance with VLOCK/VM! The key to improving your lockfile performance.*

BlueLine™
SOFTWARE INC.

New Advances in Performance, Productivity, and Planning.

ICP MILLION DOLLAR AWARDS WINNER

1500 South LilacDrive, Suite 340, Minneapolis, Minnesota 55416
800-826-0313 • 612-542-1072 in Minnesota

CIRCLE #48 on Reader Service Card ▲

of this option. An example of a case where the option would have to be used is when a set of COBOL programs is licensed from a vendor and constantly updated by the vendor. In this case it is more advantageous to use the option consistent with the vendor.

The LIB option instructs the compiler to check libraries to resolve COPY statements. The NOLIB default saves compilation time but may only be used when no COPY statements appear in a program.

Compiler Listing Directives

The Procedure Map (PMAP) option produces an Assembler language listing of the code produced by the COBOL compiler. It is usually quite long and takes a considerable amount of compilation time to produce. It makes paper companies happy when used in a compilation that will be printed. While it is quite valuable for debugging, it should only be requested when there is a chance that it will be used for such purpose, certainly not in a first compile when many errors occur.

In simple programs, for instance programs without table handling, a Condensed Listing (CLIST) may be the preferred choice over a PMAP. The two options are mutually exclusive and the default is to produce neither. While PMAP produces a full Assembler listing, CLIST only shows the offset where the Assembler instructions begin for each verb. If a dump occurs and you have a CLIST, while you will not be able to identify the precise machine language instruction in the COBOL listing that caused the ABEND, you will be able to identify which verb caused the problem. This is done by comparing the program offset computed from the dump with the offsets in the CLIST. For programmers who can understand some Assembler code and follow a PMap, the PMAP option may be useful. For programmers who will merely want to relate the displacement of an ABEND to the line in the source program that caused it, a CLIST is appropriate.

SUPMAP informs the compiler that the PMAP, CLIST, LOAD and DECK options should be ignored if an error with a severity of 'E' (return code 12) or 'D' (return code 16) occurs. In almost all cases, neither the text deck nor procedure listing will be used if a program compiles with severe errors.

The Data Map (DMAP) option produces a data description listing showing each data field used in the program, its associated BL base cell, its displacement from the beginning of the BL cell, its size, its usage and its internally generated COBOL name that is used in the PMap. During debugging, the DMap is extremely useful. It is used to locate the data fields in a dump. Also, when used in conjunction with a PMap, it is used to determine which fields are being processed in each line of Assembler code.

SXREF and XREF respectively produce sorted and unsorted cross-references of data and procedure names. While quite useful on printed listings, it is usually a waste of time to produce a cross-reference when working on-line. References to a field can be found more easily by editing a program than by editing the listing and locating the cross-reference.

Error messages, as well as many portions of a COBOL listing (PMap, CLIST, Cross-references) and COBOL output (output from FLOW, output from STATE) contain line numbers. NUM is used when line numbers existing in columns one through six of the source code are to be used in these instances; NONUM instructs COBOL to use its own generated line numbers. When the numbers in the original source code are used (NUM), it is simple to relate messages to the source. When the compiler generated numbers are used and the program contains copy members or goes through a precompilation (for example, CICS and IDMS), you must look at a listing to relate the numbers to the source. Since the length of the source program compiled will depend on the length of the copied code, there is no specific correlation between the relative line number in the original source code and the compiler generated number.

When numbers are supplied in the source code by the programmer, the SEQ option indicates that the compiler should check for ascending order of line numbers. If lines are found that are not in ascending sequence, a message is printed in the listing. The SEQ option is a carry over from when programs were entered on punch cards. To make sure the punch cards were in order (especially after a programmer attempted to place a dropped deck back in its original statement order), the option was used. The NOSEQ option instructs the compiler not to check ascending sequence of line numbers.

Debugging Options

The STATE option outputs the line number causing the problem if an ABEND occurs. At execution time, it requires a SYSDBOUT DD statement be used for output of the error message. STATE works quite efficiently. It does not change the executable code, except for an extra call during GOBACK processing. It works by keeping a table of addresses indicating where each statement begins and ends. This table is only used when the program ABENDs. The efficiency of normal processing is only slightly affected since the table used by the STATE option takes up extra storage and therefore extra paging may occur.

The FLOW option directs the compiler to output a list of the line numbers of the procedures executed directly before an ABEND occurred. The number of procedures traced is entered with the option (following an '='). When the number of FLOW procedures is not requested, a default is used. As far as efficiency is concerned, the FLOW option is quite a different matter from the STATE option. FLOW generates a call to a COBOL subroutine each time a COBOL procedure is encountered. This adds significant overhead to the program execution. Whether or not a programmer uses the FLOW option is quite subjective; I personally do not find it useful enough to outweigh its added overhead.

COUNT instructs the compiler to list each COBOL verb, its statement number and the number of times the verb is executed. The frequency statistics can be used to ensure that all parts of a program are executed during testing and to determine heavily used parts of the program. On the negative side, COUNT requires a subroutine call each time a "count-block" is executed and increases the time of program execution. A count-block is a portion of COBOL code where either all statements or none of the statements will execute (assuming that an ABEND does not occur in the middle of the block). For instance, a number of MOVE and ADD statements in a row would all be part of one count-block; an IF or GO TO statement would end the block.

Usage of the TEST option allows a program to be executed with the IBM Interactive Debug facility. The facility lets a programmer step through a program examining various data areas during execution, enables the changing of data items and even allows the path of program execution to be altered. When TEST is used, FLOW, STATE and COUNT cannot be used. By default, all four of these options are set off during a compilation. The TEST option lengthens compilation time and yields quite inefficient code. It should only

FIGURE 2

```
                COBOL II PROGRAM COMPILED WITH FLAG(I,S) OPTION
000001      DATA DIVISION
000002      WORKING-STORAGE SECTION.
000003      01  FIELD1                      PIC X(4).
000004      01  FIELD2                      PIX 9(4).
==000004==> IGYDS1089-S < PIX > WAS INVALID.  SKIPPED TO NEXT AREA A ITEM OR LEVEL NUMBER.
000005      01  FIELD3                      PIC S9(3) COMP-3 VALUE +5.
==000004==> IGYDS1159-E  A "PICTURE" CLAUSE WAS NOT FOUND FOR ELEMENTARY ITEM < FIELD2 >.
000006      01  FIELD4                      PIC X.
000007      PROCEDURE DIVISION.
000008          MOVE FIELD1 TO FIELD5.
==000008==> IGYPS2121-S < FIELD5 > WAS NOT DEFINED AS DATA-NAME. THE STATMENT WAS DISCARDED.
000009          GOBACK.
```

be used if the program will certainly be used with the interactive debugger.

Options Affecting Program Execution

The options mentioned so far, while changing the listing format and acting as debugging aids, will not change the required output of a program's execution. To be specific, if a program was written to generate a report, the report will look the same regardless of which previously mentioned options were used. The options I will now discuss can actually change the way a program executes; arithmetic computations may yield different results; a program may ABEND because the compiler does not realize which version of the program to use. All COBOL programmers should make an effort to fully understand these options.

Few programmers fully understand the TRUNC and NOTRUNC options. Depending on how a program is coded, NOTRUNC may actually be a necessary option for the program to run properly. The COBOL compiler allocates a halfword (two bytes) for binary (COMP) fields up to four digits long and a fullword (four bytes) for fields up to nine digits long. A halfword really holds up to 32,767 while a fullword up to 2,147,483,648. While a halfword or fullword COMP field in COBOL is defined as having a limit of 9,999 or 999,999,999 respectively, the actual numbers that may be put into the fields are larger. The NOTRUNC option allows a program with COMP fields to hold up to the maximum value that may be put into a halfword or fullword rather than the maximum value determined by the PICTURE clause.

In programs that manipulate addresses (such as BLL cells in CICS programs) and other programs manipulating fullwords (like database keys in IDMS), the NOTRUNC option may be required. Regardless of the application, the NOTRUNC option enables the compiler to generate substantially more efficient code when manipulating binary (COMP) fields.

The TRUNC option ensures that after arithmetic is performed on a COMP field, if the value of the field exceeds the maximum value of the field definition, the result will be truncated to the PICTURE size. Extra machine code is added to accomplish this. It is rare that a program wishes to ensure that a binary (COMP) field exceeding its defined number of digits be truncated to the defined size. Except in these rare cases, the NOTRUNC option should always be used. Every company should leave the NOTRUNC option as the default for all compiles.

The DYNAM option makes all CALLs in a program into dynamic calls. This allows a called program to be loaded at the time of its invocation rather than forcing it to be linkedited into the same load module as the calling program. While CALLs to a program whose name is defined as a data field rather than a constant in a CALL statement are automatically considered dynamic, calls that have the program name coded as a literal are static unless the DYNAM option is used. CICS programs require NODYNAM.

Be careful not to call a program dynamically from one program and the same program statically from another. This can result in two copies of a program being placed into storage and may result in an ABEND or other unexpected outcome.

The RESIDENT option makes calls to the COBOL library routines dynamic. NORESIDENT means that calls to the library routines will be static, requiring the routines to be made part of the program load module. The RESIDENT option can be used without the DYNAM option. However, when the DYNAM option is chosen, the RESIDENT option is automatically in effect.

ADV tells the compiler that when the WRITE . . . ADVANCING clause is used, the carriage control character (first byte of the record) will not be specified in the program but is to be added automatically by COBOL. NOADV says that the carriage control byte is defined in the program's record definition. If a program was writing out 132-byte records and the ADV option is used, the output record would be defined as PIC X(132); with NOADV it would be defined as PIC X(133) and the first byte would be left unused in the program.

OPT instructs the compiler to perform code optimization. The compilation will take longer if this function is requested. Although the optimization does make the program process faster, non-IBM products such as CA-Optimizer from Computer Associates (Garden City, NY) do a much better job. However, CA-Optimizer has to run as a separate step and takes much longer than using the OPT option of the COBOL compiler.

COBOL II Changes And Enhancements

A number of options have been changed or updated. (For a chart of options that had their names changed see "COBOL II: Its Differences and Idiosyncrasies" in the May 1989 issue of *MAINFRAME JOURNAL.)*

FLAG (x,y) is an enhancement of the FLAGW and FLAGE options of the VS COBOL compiler. The FLAGW and FLAGE options respectively inform the compiler to only list error messages with a warning level ('W') or higher or with an error level ('E') or higher. The 'x' in the FLAG option serves a similar function and tells the compiler to only list error messages with a severity greater than or equal to the value substituted for 'x'. The 'y' value is a new feature that requests errors with a severity greater than or equal to 'y' ('y' cannot be less than 'x') to be placed in the listing at the point that they are detected (see Figure 2.) This is usually the point at which they occur but due to compilation logic may be somewhat later. The error messages are also placed at the end of the listing when this option is used.

When the TEST option is used in COBOL II, either the batch mode debug-

See COBOL *page 93*

DOWNTIME. WITH AN EXCUSE.

Ordinary diagnostic tools do an okay job of telling you what went wrong after your system goes down.

There's only one problem.

Your system is down.

NO DOWNTIME.

STABILIZE®/CICS keeps your system up and running.

STABILIZE automatically ■ detects ■ diagnoses AND has the intelligence to ■ **DYNAMICALLY REPAIR CICS** to prevent system failures.

And, with each new release of STABILIZE, we update its knowledge base so that it keeps learning how to detect, diagnose, and repair more and more CICS problems thus preventing CICS outages.

Your system stays up.

The choice is yours, but unless you choose STABILIZE/CICS, there's no choice at all.

Limited time offer: License STABILIZE by August 25, 1989, and receive, at no extra cost, the new "CENTRAL" option for automating CICS operations with central message control.

My choice is NO DOWNTIME.

☐ Send me a free STABILIZE Presentation Diskette.
☐ I want to trial STABILIZE for 30 days. Call me.
This request qualifies me to receive, at no extra cost, the CENTRAL option when I license STABILIZE by 8/25/89.

Name/Title

Company Name

Address

City/State/Zip

Mail Stop Telephone

Mail to: On-Line Software, Two Executive Drive, Fort Lee, NJ 07024 or Call

800-642-0177

201-592-0009

On-Line Software INTERNATIONAL

The Safe Buy.

All our products are offered with a life-time trade-in guarantee so that the money you spend today is always available to meet your changing needs tomorrow.

On-Line Software offers consulting, education and software products. We specialize in CICS, DB2, and CASE technologies for application development.

MJZDG9

CICS RADAR® is a registered trademark of Compuware Corp.
Eyewitness® is a registered trademark of Landmark Systems Corp.

DB2 Storage Groups
A Viable Production Alternative

By Barry Lewis

One of the main objectives of a relational database is to allow data users to get what they want without having to know many highly technical, computer oriented details. Hopefully, the DBMS can anticipate the technical needs of the user and perform those functions automatically. However, this type of automation is only useful if the software is smart enough to make acceptable decisions with regard to performance and resource usage.

In a number of ways, the DB2 storage group concept as supported in Version 1 failed to make those decisions well enough to be utilized in a production environment. Some installations utilize storage groups on their test systems, but few utilize them in production. Version 2 offers enhancements that resolve some of the problems that have caused users to shy away from storage groups.

For those readers who are not familiar with the purpose of a storage group, allow me to give a brief description. DB2 relies totally on VSAM for allocating and cataloging disk space (entry sequenced datasets in Version 1, linear datasets in Version 2) where table data is to be stored. As with all VSAM datasets, Access Method Services (AMS) must be invoked to define those datasets. The question arises as to who performs this definition; the user with his handy-dandy IDCAMS job or DB2 through some automated process.

IBM offered a choice. The user could define the datasets and then supply the dataset name to DB2 indirectly through the USING VCAT clause or DB2 could be requested to perform this function automatically as datasets were needed. The latter was accomplished by inventing the storage group, which is actually nothing more than a list of volumes where DB2 has permission to put datasets. When the user issues a SQL CREATE statement for either a tablespace or an index, he/she specifies the storage group name (in the USING STOGROUP clause) and DB2 goes to that list of volumes, finds one with adequate space and invokes AMS to allocate the dataset. If the tablespace or index is later dropped, DB2 sees that the associated VSAM dataset is deleted as well.

Strategically, the concept of storage groups fits well with IBM's DASD management strategy for MVS/ESA. It allows datasets to be assigned to profiles of definition parameters called storage classes and management classes, thus simplifying the JCL required by the user in defining the datasets.

The purpose, then, of a storage group, is to eliminate the hassle of defining and deleting datasets. However, two major problems existed.

Problem One

Suppose the storage group STOGRP01 represented three volumes as shown in Figure 1. When directed to this STOGROUP, DB2 always places the dataset on the first volume in the list with adequate space. The second (and any subsequent) volume is not used until the first fills up. This uneven distribution is generally not acceptable from a DASD performance standpoint. One suggested solution was to create a STOGROUP for each volume. Then the desired volume could be designated by supplying the associated STOGROUP name. This works okay for small to medium tablespaces but not so well for large tablespaces that may need to span across volumes. A tablespace is owned by a STOGROUP, meaning that its dataset(s) can only exist on the volume list for the STOGROUP. If the STOGROUP has only one volume, the tablespace must exist on that one volume. Some relief was offered in Version 1 with the *partitioned* tablespace, a special implementation allowing the tablespace to be broken up into partitions based on ranges of key values. Each partition utilizes a separate dataset and can be directed to a different STOGROUP. A configuration as shown in Figure 2 would distribute the data evenly across volumes.

Problem Two

Probably a much more severe limitation of Version 1 STOGROUPs was the inability to change the size of the dataset without dropping and recreating the tablespace or index. This is a common administrative task to alleviate multiple extents or out-of-space conditions. The problem was due to the fact that the ALTER TABLESPACE statement did not include the ability to change the values of PRIQTY and SECQTY. Other problems

FIGURE 1

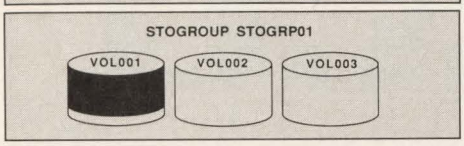

FIGURE 2

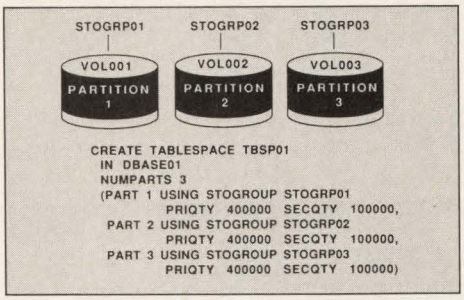

existed due to lack of functionality in the ALTER TABLESPACE statement. The following tasks could not be performed without unloading the data, dropping and recreating the tablespace and reloading the data:

- Move the tablespace to a different STOGROUP (and possibly a different volume)
- Convert a user managed tablespace to a DB2 managed tablespace
- Convert a DB2 managed tablespace to a user managed tablespace.

DB2 Version 2 offers considerably more flexibility with regard to altering tablespace (and indexes). The following examples should demonstrate these new facilities. Although the examples involve a tablespace, the enhancements apply to indexes as well (through the ALTER INDEX statement).

Example One: Altering Space Allocations

Suppose tablespace DBASE01.TBSP01 has run out of space. The current allocations are PRIQTY = 40000 and SECQTY = 10000:

```
ALTER TABLESPACE DBASE01.TBSP01
      PRIQTY 80000
      SECQTY 20000 ;
```

This change does not take effect immediately. Either REORG or RECOVER will cause the dataset to be deleted and redefined using the new allocations.

Example Two: Moving A Tablespace To A Different STOGROUP

First note that a volume can be owned by more than one STOGROUP as indicated in Figure 3. This means that moving to a different STOGROUP may or may not involve a volume change.

Assume DBASE01.TBSP01 currently belongs to STOGRP02 and resides on VOL002 and you want to move it to STOGRP01.

```
ALTER TABLESPACE DBASE01.TBSP01
      USING STOGROUP STOGRP01 ;
```

Since STOGRP01 also owns VOL002, no volume change is necessary. However, since VOL002 is not the first volume in the STOGRP01 list, the next REORG or RECOVER may move the dataset to VOL001 if adequate space is available there.

To later move the same tablespace to STOGRP03, you would have to actually move the dataset, because STOGRP03 does not own VOL002. To perform this task:

1) Stop tablespace DBASE01.TBSP01
2) ALTER TABLESPACE DBASE01.TBSP01
   USING STOGROUP STOGRP03 ;
3) Restart tablespace DBASE01.TBSP01 as utility only (UT)
4) Execute the RECOVER or REORG utility
5) Restart tablespace DBASE01.TBSP01 as READ/WRITE (RW)

If PRIQTY and SECQTY are not specified in the ALTER statement, the new dataset created by RECOVER or REORG will be allocated with the same attributes as the old one. As is the case with any altered tablespace parameters, the changes do not take effect until REORG or RECOVER when STOGROUPs are being used.

Example Three

Example three involves converting a user managed tablespace (USING VCAT) to a DB2 managed tablespace (USING STOGROUP).

Assume the tablespace was originally created as follows:

```
CREATE TABLESPACE TBSP01
      IN DBASE01
      USING VCAT USER20 ;
```

The USING VCAT node informs DB2 that a user-defined dataset already exists named:

```
USER20.DSNDBC.DBASE01.TBSP01.I0001.A001
```

The only part of the name that users have the flexibility to specify is the first qualifier USER20 that was specified in the USING VCAT clause.

Now assume that the dataset currently resides on VOL002 and you want to convert the tablespace to be DB2 managed through STOGRP02 (see Figure 3).

```
ALTER TABLESPACE DBASE01.TBSP01
      USING STOGROUP STOGRP02
            PRIQTY 80000
            SECQTY 20000 ;
```

The dataset will be used as is until RE-

FIGURE 3

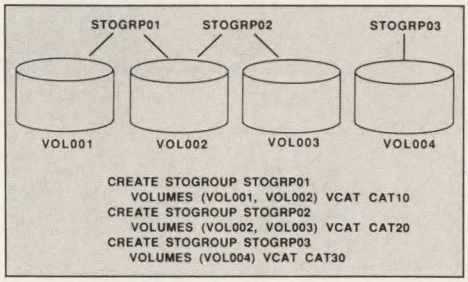

COVERY or REORG. At that time, DB2 will delete the USER20 dataset, find space on STOGRP02 and allocate a new dataset named:

```
CAT20.DSNDBC.DBASE01.TBSP01.I0001.A001
```

with PRIQTY = 80000 and SECQTY = 20000. From that time on, the dataset will be managed by DB2. Note that in this example, specification of PRIQTY and SECQTY are required; since the dataset was previously user managed, DB2 has no defaults.

Note that none of the examples required dropping and recreating the tablespace. These additional functions should eliminate many of the problems noted previously. Although Version 2 offers no additional help in solving the DASD management problem, this can be overcome by creating a STOGROUP for each volume and using partitioned tablespaces to distribute data for large tablespaces across volumes as desired.

There are several products on the market by third-party vendors that would also automate the process of changing data structures and resource requirements. These provide greater human control with minimal additional effort.

Still, regardless of what utilities and procedures are used by the DBA, the combination of the new features of DB2 Version 2 and the use of partitioned tablespaces may allow the storage group to become a viable alternative for production systems.

ABOUT THE AUTHOR

Barry Lewis is the DB2 curriculum manager for a Dallas-based computer education company. He has written numerous DB2 courses and conducts seminars on a regular basis at major corporations across the country. (214) 252-5354.

Succeed

PLATINUM *technology, inc.*
The DB2 Company™

PLATINUM *technology, inc.* is the only full service software company specializing exclusively in DB2. All our effort and energy is focused on delivering the broadest array of quality solutions for the DB2 user. As a result, IBM® has designated PLATINUM a Business Partner through the Authorized Application Specialist program specifically for DB2.

PLATINUM PRODUCTS

PLATINUM offers a complete line of software tools, education, and published products to ensure your success with DB2.

Software Products

The PLATINUM product portfolio consists of a complete family of administration, development, and end user DB2 software products. All are compatible with DB2 V1.3 & V2.1. The software products include:

- **RC/Query™** - A comprehensive DB2 catalog query tool.
- **RC/Update™** - The industry's best DB2 object management and data editing tool.
- **RC/Migrator™** - A complete object and data migration tool.
- **RC/Secure™** - An extensive DB2 security management tool.
- **PLATINUM Database Analyzer™** - The DB2 database and DASD analysis, audit, and management tool.
- **PLATINUM Report Facility™** - The DB2 query and reporting system for developers and end users operating in TSO and CICS environments.

DB2 Education Courses

A complete series of hands-on DB2 training courses. PLATINUM courses cover all aspects of DB2, QMF, and CSP. All courses are available either at your facility or at our Corporate Education Center.

- Introduction to DB2
- DB2/SQL Application Programming
- DB2 Application Planning and Database Design
- DB2 Database and System Administration
- Using DB2 and QMF
- CSP Application Programming

Published Products

The most recognized and authoritative DB2 standards, methods, and guidelines for DB2 implementation.

- **PLATINUM DB2 Guide/Online™** - The industry's leading standards manual for design, development, and administration of DB2 systems.
- **The PLATINUM Reference™ for DB2** - The quick, pocket-sized reference for DB2 information.

Support

All products and services carry our PLATINUM Quality Assurance Guarantee. Support is available 24 hours a day, 7 days a week via our toll-free hot line.

WORLDWIDE AVAILABILITY

PLATINUM's products and services are available around the world through our Affiliate Network. PLATINUM's full service capabilities include local support, education, and superior DB2 professional services.

with DB2!

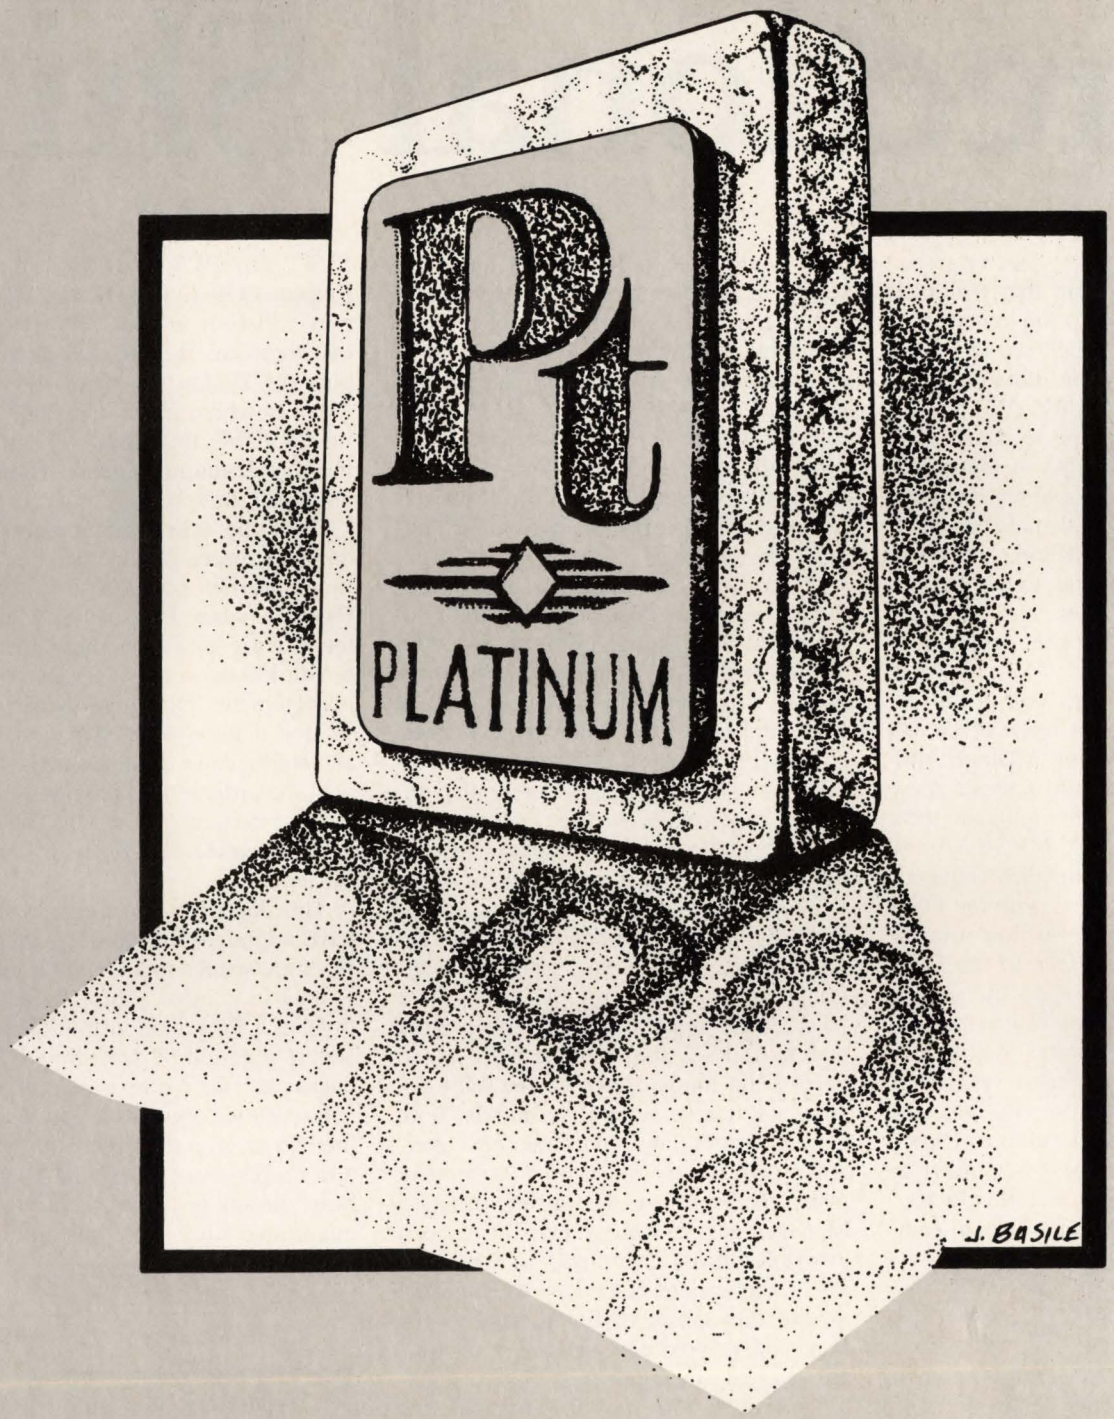

SAA–IBM's Directional Shift

By Ira W. David

As time passes, it is becoming increasingly evident that the focus of IBM's push to implement Systems Application Architecture (SAA) is gradually shifting. While initially stressing portability, uniformity and consistency, more ink today goes to cooperative processing and its benefits in the corporate *enterprise*. This is not to say that there have been retractions or denials but rather that the emphasis is different. In some cases IBM may be holding to the letter of an announcement, rather than what the user base initially read as the meaning of the announcement.

As a simple example, look at the announced implementation plans of Dialog Manager and Presentation Manager. Initially, the direction was to provide these tools under OS/2 and then to migrate them to operate under OS/400, MVS/ESA and VM/XA/SP. However, this has changed.

The direction still starts with the OS/2 version, but the next step is now a mainframe *cooperative interface* to the OS/2 facility. This would provide the Dialog Manager and Presentation Manager *functionality* on the mainframe. Whether or not this complied with earlier statements is a matter of semantics. (In fact, the pieces are no longer referred to as Dialog and Presentation Managers but rather as Dialog and Presentation Interfaces.) Whether a true mainframe implementation would ever exist was left open for future re-evaluation. Furthermore, there was a glaring lack of mention of the OS/400 implementation.

To those users with the appropriate hardware, software and network connectivity, this may be an improved approach, off-loading cycles to the less-expensive Intelligent Workstation (IWS). However, for most of the world, this is not an equivalent alternative.

Reasonable estimates for a *cooperatively-configured* IWS with an SAA configuration run from $5,000 to $10,000. To implement mainframe applications reliant on such end-user devices would mean mass replacement of Mainframe Interactive (MFI) or CRT devices as well as DOS-only PCs. This is a potentially immense expense. Most users cannot, and will not, make such a changeover without major analysis and justification.

To further complicate matters, shops typically run a mixed bag of end-user devices: some OS/2 units, some MFI terminals and some PCs running DOS. With IBM's new approach, everything other than those PS/2 units running OS/2 would have to be lumped in with the MFI devices and treated as dumb terminals. With earlier assumptions and functions operating totally on the mainframe, there were many pieces of the Dialog/Presentation Manager puzzle which could have been available across the board to MFI and IWS. Applications could have been geared toward this subset of functionality. This is no longer possible with the new strategy. For the average user, this is a major consideration; for the Independent Software Vendor (ISV) it can be earth-shattering.

Furthermore, it is the responsibility of the application program to decide whether to invoke the Dialog Interface. This may really be a decision of whether the end device is supported. In the earlier version, the type of connection would be the giveaway with LU2 being the key for MFI and LU6.2 for IWS. However, the PC can support LU6.2. In fact, the PC can support CPI-C if anyone wants to code up the interface. But, the PC does not support the Dialog Interface. So you must look elsewhere for an indicator. (According to IBM, the Cross System Product (CSP) will automate the decision, but you might ask how.)

Even the schematic representation of how the pieces fit together has changed. In the earlier portrayals, it looked like Figure 1, where Presentation Manager was the driving piece handling the physical screen I/O as well as the task/dispatching management. Dialog Manager was considered a *user* of Presentation Manager, an intermediate application so to speak, with the programmer having the option of either interfacing through Dialog Manager or dealing directly with Presentation Manager.

This description of the pieces seems to have migrated to Figure 2, where the application speaks to a Dialog Interface that is now in control. This Dialog Interface invokes a Presentation Interface if graphics are required or deals directly with the display device for text-only.

Given the cooperative nature of today's implementation direction, this new structure makes sense. Only the Dialog Interface portion has to have a stub on the mainframe since the application program

FIGURE 1

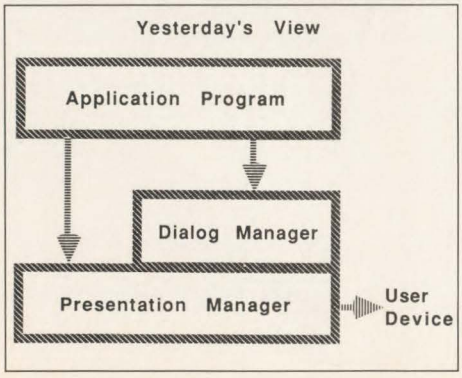

FIGURE 2

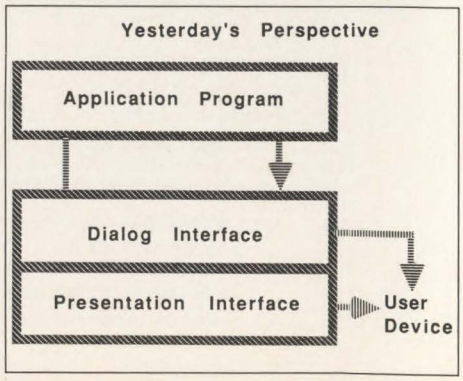

cannot speak directly to the Presentation Interface. The Presentation Interface handles the graphics screen driving, if any, and the Dialog Interface handles the rest; that is, the textual screen handling as well as the low-level processing control on the IWS and the higher-level processing on the mainframe (on the *other* side of the cooperative link).

This is logical for the end-user configuration being recommended by IBM. However, it leaves more than the MFI in the lurch. That group is now joined by all of the micro users (PC and PS/2) who cannot, will not run or are not running OS/2. Until such time as an ISV comes up with a Dialog/Presentation Manager or Dialog/Presentation Interface look-alike to run under DOS, this group is no better off than users of the basic MFI. Graphics are possible, but no more so than on a simple graphics terminal (assuming CGA or higher display) and even then only with a direct application invocation of GDDM.

Why would IBM take a stance like this? After the claims of support and extensions for the CRT user, why is this group being pushed aside in the Dialog/Presentation Interface plans? I can think of two potential reasons, both based on extrapolation and assumption. These are personal opinions only, from what I have seen and interpreted. Clearly IBM is not letting me in on their secrets.

On the one hand, to say that user acceptance of OS/2 has not been what IBM would have hoped for would be to grossly understate the obvious. This new scenario does give some additional reason to push users into the additional expenditure. While IBM appears confident that the OS/2 strategy will pay off in the long run, this redirection may have some short-term marketing benefits.

Personally, however, I think that there is an even stronger possibility. The myriad of pieces announced in the SAA unveiling of March 1987 was staggering and was added to at an alarming rate with things like CPI-C, DDM, new languages to include and more. Technical resources for so many major parallel efforts may be proving too much to handle, even for a giant like IBM, especially in light of some of their more recent financial reports and staffing cutbacks.

Everything has its limits and IBM's user base has been screaming about the lack of short-term deliverables. We have seen a number of the expected announcements and deliveries delayed or reduced. The initial (OS/2) Presentation Manager was less robust than hoped for. The infamous repository is still amorphous. SAA and SAA-related announcements have been slowing to more closely resemble reality.

Clearly, the cooperative implementation is easier than a complete version. It appears to fulfill the letter of the original announcement, even if not the assumed spirit. Also, the complete mainframe implementation has not been precluded for the future, merely for the present. In other words, it would appear that IBM may have "changed the specs" to better reflect what could be delivered today.

ABOUT THE AUTHOR

Ira W. David is Vice President of MIS at Idelman Telemarketing, Inc. in Omaha, NE and was formerly Executive Vice President of On-Line Software International in charge of SAA impact analysis.

Selecting Artificial Intelligence Software

By Howard W. Miller

Artificial intelligence is a broad concept encompassing a number of different disciplines, including cognitive psychology, decision theory, operations research, machine learning, robotics, natural language processing and expert system technology. Artificial intelligence involves the creation of computer software that emulates the way people solve problems or cloning the way humans think. Business professionals are most likely to come in contact with the discipline of expert systems; therefore, for purposes of this discussion, artificial intelligence is equated to expert systems.

Experts solve difficult problems, explain the solution, learn from the problem solving process, explain the relevance of the solution and, maybe most importantly, experts are capable of realizing when they do not know something. An expert system, like a human expert, gives advice by requesting information specific to the problem under consideration and by drawing upon its store of knowledge.

The incentive behind expert systems is contained in the attributes of knowledge. Knowledge is perishable and its longevity is tied to the expert. Expert knowledge is scarce and difficult to accumulate, pass on and utilize. Further, expert knowledge is often vague, inconsistent and widely dispersed over many widely distributed experts. In response to these attributes of knowledge, expert systems seek to preserve, clone and apply knowledge. Expert systems seek to pass on knowledge to another generation of experts or users and to encourage its growth and expansion. Expert systems attempt to make the knowledge more precise and systematic and to collect it into a knowledge base.

Expert systems, therefore, provide the benefits of making knowledge more readily available. The expert system is impartial in its decision making and it has total recall. It provides the opportunity to share knowledge over a large user base and it expedites routine decisions. Finally, expert systems conserve valuable experience in an organization and can act as a tutor to pass it on to trainees.

The history of expert systems is as old as information technology. Early computer practitioners had high expectations for artificial intelligence. These early knowledge engineers attempted to replicate the problem-solving capabilities of human experts, but the expectations did not materialize as quickly as anticipated. The process of human thinking is far more complex, less structured and more elusive than knowledge engineers assumed. Furthermore, the type and intensity of computer processing exceeded the capabilities of available computers.

Early expert systems required immense amounts of human labor and computing, using special purpose computers that were not readily available. From a business perspective, expert systems were impractical; both the human involvement and the computing resources were more expensive than practical. However, investments in expert system technology has risen significantly over the last few years. A September 1986 article in *Modern Office Technology* reported that expenditures on expert systems increased from a reported $13 million in 1983 to a projected $800 million in 1987. In light of these investments in expert systems, it is obvious that the attitude of business is much more receptive toward expert systems.

Uses of Expert Systems

Application of Expert Systems

There are a number of different ways to look at expert systems. Expert systems can be categorized by their target markets (medicine, industry, geological exploration, insurance or military) or by their hardware environment (micro, mid-range, mainframe or AI workstation). However, a more common way to categorize expert systems is by the target user market: non-expert, technician or expert.

Advise Non-experts

These expert systems supply expert knowledge directly to non-experts. These expert systems provide advice on taxes or tax preparation, supply financial planning expertise, perform investment analysis and provide advice in areas such as medicine, gardening and hobbies.

Improve Performance of Technicians

These expert systems improve the performance of the technician. The systems improve performance through such activities as identifying the need for preventive maintenance, configuring equipment, performing diagnostics and assisting with the operation of complex equipment.

Aid or Outperform Experts

These expert systems aid or even outperform the experts. The systems aid experts through such activities as seismic analysis, medical diagnosis, oil prospecting, financial analysis and tax planning. Other systems outperform the experts in such areas as process control, real-time financial trading and playing chess.

Expert System Examples

The following are twenty-five examples of expert systems that have been developed and are in use at major corporations or are commercially available for general use.

1. **AUDITOR:** Assists a corporation in analyzing the allowance for bad debts and accounts receivable.
2. **AUTHORIZER'S ASSISTANT:** Performs credit authorization searches at American Express and makes recommendations to the authorizing agent.
3. **CASH VALUE:** Assists in capital project planning: advises on NPV, cash flow, payback and risk analysis.
4. **CONSULTANT:** An expert system that helps IBM field service representatives prepare price bids.
5. **CORP-TAX:** Assists accountants with Section 302(b) redemptions.
6. **DENDRAL:** Elucidates chemical structures from mass-spectral data.
7. **DRILLING ADVISOR:** Diagnoses, solves and helps avoid problems with oil drilling rigs.
8. **EDP AUDITOR:** Aids auditors in assessing advanced electronic data processing systems.
9. **EXPERTAX:** Helps accountants at Coopers & Lybrand to review ways their clients can accrue taxes and assists in providing tax planning advice.
10. **FINANCIAL ADVISOR:** Gives advice on projects, products, mergers and acquisitions as if conversing with a senior financial consultant.
11. **GUIDON:** Performs medical teaching.
12. **HASP:** Understands complex, noisy, analog signals for such things as submarine detection and identification.
13. **MACSYMA:** Performs mathematical manipulations such as integration, differentiation and simulation equations.
14. **MCYIN:** Performs medical diagnosis and therapy recommendations.
15. **MUDMAN:** Analyzes the drilling fluids or "muds" that are pumped down the shaft to facilitate drilling by lubricating.
16. **PDS:** A Westinghouse expert system designed to monitor steam turbines and to make maintenance recommendations.
17. **PLAN POWER:** Takes into consideration a financial situation and then matches needs with the most appropriate financial products and services.
18. **PLATINUM LABEL:** A general accounting expert system that includes seven expert packages: accounts receivable, accounts payable, general ledger, sales order, inventory, sales analysis and guide database kit.
19. **PROSPECTOR:** A mineral prospector that identifies sites for ore deposits.
20. **R1:** Designs complex computer configurations.
21. **SACON:** Advises on the usage of complex software products.
22. **TAX ADVISOR:** Provides tax advice to help clients arrange financial affairs to minimize income and death benefits taxes.
23. **TAXMAN:** Evaluates the consequences of proposed business reorganizations.
24. **TICOM:** An expert system for modeling and evaluating internal financial controls.
25. **XCON:** An expert configuration tool, designed by Digital Equipment Co. (Boston, MA), used to check sales orders and design the layout of each order analyzed.

The Knowledge Engineer

The array of expert system application opportunities is diverse and the opportunities are almost limitless. However, isolating these opportunities requires an awareness on the part of the organization that such applications are possible, that the results are valuable and that there will be a commitment to use the results. Further, the technology is new to business and it requires a long-range investment in knowledge engineers and expert system software tools.

Knowledge engineers are themselves experts. They specialize in isolating information from experts and in understanding the strengths and weaknesses of their chosen expert system tools. A knowledge engineer typically identifies an application suitable for solution with an expert system, isolates information from an expert, develops a prototype of the expert system and in close cooperation with the expert, develops a working expert system. These systems are then integrated into existing automated systems and are turned over to the expert and the system user for subsequent support.

The conventional software engineer also works with experts; however, the relationship between the knowledge engineer and the expert is much more intense. Conventional software engineers are concerned with information or business flow. They work with an expert to isolate this business information flow and afterwards design software systems to automate this business information flow. The knowledge engineer, however, is concerned with the thought process of the expert. The interaction between the expert and the knowledge engineer continues from information isolation through developing the working system. Frequently, the expert takes over the subsequent support for the expert system.

Based on the size of the project at hand, more than one knowledge engineer may be required. However, most commercial applications use a single knowledge engineer and one or more conventional software engineers. Most early expert systems were programmed from scratch using LISP or PROLOG, but today, the most common media is an expert system shell or a programming environment or toolkit.

The use of tools such as programming environments and expert system shells has improved the productivity of the knowledge engineer and has made expert systems much more commercially accessible. However, the success of an expert system still largely depends on the ability of a knowledge engineer to isolate expert information and to use a programming environment. Isolating expert knowledge continues to be a labor-intense process for both the knowledge engineer and the expert.

Types of Expert System Development Software

As stated above there are three principal alternatives for developing expert systems: an expert system shell software environment, a more general purpose programming language and a programming language toolkit. The primary difference is in the amount of knowledge engineering experience required and in the amount of effort required to develop the expert system. The alternatives for developing expert system software range from the shell requiring the least amount of programming effort to the programming language requiring the most.

Expert System Shell

An expert system shell is so called because it is empty of any knowledge. Shells usually consist of four components: a knowledge base, an inference engine, a user interface and a knowledge encoding facility. It may or may not include an interface to a traditional hierarchical or relational database.

See Artificial Intelligence *page 78*

VSAM Dataset Hazard

By Howard Glastetter

DASD performance can be most improved by the elimination of unnecessary I/O. (*NO* I/O is the fastest you can do.) This short article concentrates on a single DASD VSAM performance issue: the effect the Control Area (CA) size has on the VSAM index structure and I/O performance. A CA is a variable amount of DASD space that VSAM uses to hold data records. All VSAM data and index space is made up of one or more CAs.

CA Size Affects Index Structure

CA size influences the VSAM index structure. A VSAM Keyed Sequential Dataset (KSDS) runs most efficiently if the index has as few levels as possible. Overhead increases (often geometrically) as the index levels increase. This is because an index record from each level must be read to locate desired data. Index levels increase when a level cannot point to all the data or all the index records in a lower level. One level gives optimum performance, but this can only occur if

Is The VSAM IMBED Parameter Passé?

Almost every VSAM manual says the IMBED parameter improves performance. It places the SSI (VSAM's lowest level index) on the top track of the CA that it indexes, then replicates the pointers around the track. This reduces arm movement and rotational delay when accessing an SSI just prior to accessing the data in the CA pointed to by that SSI.

Cache Issues

However, if your dataset is on a DASD volume that uses a CACHE memory controller, this feature will actually *degrade performance*. Cache controllers stage a track of data at a time to their memory. Replicated data will waste 90 percent plus of the track that could be filled with more of the index. The next read of the index could be filled from high-speed cache memory rather than DASD.

Place *heavily read* VSAM datasets on volumes that use cache. *Do not use IMBED*. Ideal candidates have several reads per track and little (five percent maximum) prime shift write or update activity. Separating the index component to a CACHEd volume can also be effective even for datasets that receive a fair amount of update activity. (Separating the index component to a different volume can also be an alternative to IMBED in a non-cache environment.)

When caching an entire dataset, here are a couple of tips. Small random datasets usually do better than large ones. If the dataset is read only, then use 0 percent CI FREESPACE and a CI size of 4096. This allows the maximum amount of data to be staged to cache memory, increasing chances of *read hits* from future accesses to the dataset.

Other IMBED Concerns

IMBED consumes space. It uses a track of each CA. On a 3350, this was one of 30 tracks (3.3 percent). On a 3380, this is one of 15 tracks (6.7 percent). 3380 tracks are 2.5 times larger than 3350s. A 4K SSI replicated four times on a 3350 will have 10 copies on a 3380. All IMBEDed SSI pointers in a CA must be changed when a CI or CA split occurs. This means overhead/risk. Partial updates of IMBEDed indexes prior to a system crash can destroy a VSAM file. Granted, the risk is small, but it is there.

Note too, when your IMBEDed file goes to extents, SSI indexes are loaded out to each extent (index separation does not occur with NOIMBED). Fragmented indexes without good buffering cause *degraded performance*. Since IMBED increases file size, it is often part of the reason for extents.

VSAM defaults to only enough BUFFERSPACE to hold one Index CI in memory at a time. Unless the BUFFERSPACE is purposely made large enough to hold the high level portion(s) of the index in memory, DASD arm movement could actually be increased by using IMBED. (In

the dataset occupies a single CA. The largest a CA can be is one cylinder (15 3380 DASD tracks). The smallest is one track. Many active VSAM datasets occupy less than one cylinder, yet have two index levels and much more overhead than necessary. Why?

You cannot ask for a specific CA size; VSAM assigns you one based on file allocation. If you allocate your data component as CYLINDERS (1 1), the assigned CA size will be one cylinder. Great, that is what you want. Your file gets one index level with one index Control Interval (CI) as long as it stays less than one cylinder. A CI is normally a physical block of DASD space that holds index or data records in a CA.

However, if that same file is defined as TRACKS (15 1) then the assigned CA size will be one track. You will get two levels of indexing if your file size exceeds a track. You will get as many secondary indexes as you have tracks of data.

What causes VSAM to make this inefficient assignment? When you allocate a data component in tracks, VSAM creates a CA size that fits the smallest of the primary or secondary space allocations, which is a track in the above example. Therefore, allocate your small VSAM files in cylinders or make secondary and primary allocations equal; that is, TRACKS (15 15), if you are using tracks. (This article mainly discusses small VSAM files of one cylinder or less. VSAM files larger than one cylinder should always be allocated in CYLINDERS. Some large VSAM files that should have two index levels have three because of track allocation.)

CA Split Concern

A small CA size can cause another problem besides unnecessary index levels. Records added to small CA sizes have a much higher chance of CA splits (an amoebic reaction when one CA becomes two in order to hold the data). A CA split can cause more than 100 times the overhead of a normal record addition and waste up to 50 percent of your disk space.

IMBED Caution

Additionally, if you use the IMBED parameter, you further lower the CA size by order to get to the next SSI CI, the *arm* would have to move to read the high level index then back to the read SSI.) If you ask the applications programmer to use IMBED, (s)he must give thought to overriding the inefficient VSAM BUFFERSPACE default. Will that happen? Not always. With some people — not at all. Another alternative is to write an exit or buy a commercial package to do creative buffering at open or allocation time.

Some VSAM "experts" are lately recommending not using IMBED in a 3380 environment. IMBED may help files that are large, need good performance and do not benefit from cache. However, there is a price and IMBED should only be used after much thought — *if at all*.

Is IMBED passe? The answer is, yes, if you are using 3380 DASD, if you cannot automatically override BUFFERSPACE defaults and, most importantly, if cache controllers are accessing your VSAM. *H. G.*

a track to hold a replicated Sequence Set Index (SSI). An SSI is the lowest level index VSAM uses. There is one SSI for each CA if there is more than one CA. IMBED will also add to any CA split overhead (all replicated SSI pointers of the old and new CA will have to be modified). IMBED is of negative value for any dataset less than a cylinder. It also degrades performance in the 3880 DASD cache controller environment. (Why stage 10 or more *replicated* SSIs to cache controller memory when you can stage 10 or more *different* SSIs instead?)

Prudent Use Of Freespace

The VSAM FREESPACE parameter also influences the CA. FREESPACE (10 20) tells VSAM that you want to reserve 10 percent of each CI and 20 percent of each CA. FREESPACE is good if your file receives a lot of record additions. CA FREESPACE, especially, will reduce CA split activity. Be aware that all actively updated VSAM datasets should also be reorganized often to avoid CA splits. Also be aware that FREESPACE is useless or even harmful if the dataset is never or rarely updated. A FREESPACE (10 20) issued to a static file that would barely fit into one CA will be forced into two CAs with two index levels and less data in each CI.

Real-Life Example

Correction of CA related problems can result in dramatic performance improvements. Recently, a customer at our site had a small VSAM dataset that received *half a million* physical reads of the index from a single batch job. At first look, this seemed to be a BUFFERSPACE problem. VSAM defaults to only enough BUFFERSPACE to hold one index CI in memory at one time. However, the real problem was due to poor CA size, IMBEDed SSIs and unnecessary FREESPACE. When all the data was forced into a single CA, the physical reads of the index dropped to one. That is a 99.9998 percent reduction in I/O. The job cost also dropped by several hundred dollars and completed in much less than the previous *normal* time.

Conclusion

- If your file occupies more than 10 tracks, *allocate in cylinders*
- If you allocate in tracks, *make secondary and primary equal*
- DO NOT USE IMBED
- If the dataset will not be updated, *use 0 percent CI and CA FREESPACE*
- If there will be a lot of record inserts, use generous CA FREESPACE and *reorganize frequently.*

If you have ideas, comments or suggestions, write or call:

Bob Thomas, editor-in-chief,

MAINFRAME JOURNAL, *PO Box 551628, Dallas, TX 75355-1628, (214) 343-3717.*

ABOUT THE AUTHOR

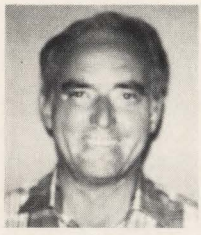

Howard Glastetter is with the State of Washington Department of Information Services in the computer performance area. His entire data processing background has been with the State of Washington. It includes 20 years of IBM OS/MVS mainframe experience with fifteen years in applications. The last five years he has worked in the areas of computer performance, capacity planning and DASD performance. State of Washington, 12th and Franklin, Bldg. 2, MS OB-01, Olympia, WA 98504, (206) 586-1019.

JOBTRAC™

Comprehensive Production Control Software that includes:

- Job Scheduling
- Job Restart and Recovery
- Sysout Archival and Online Retrieval
- Console Message Recognition and Reply
- System Operator Command Scheduling
- And More...

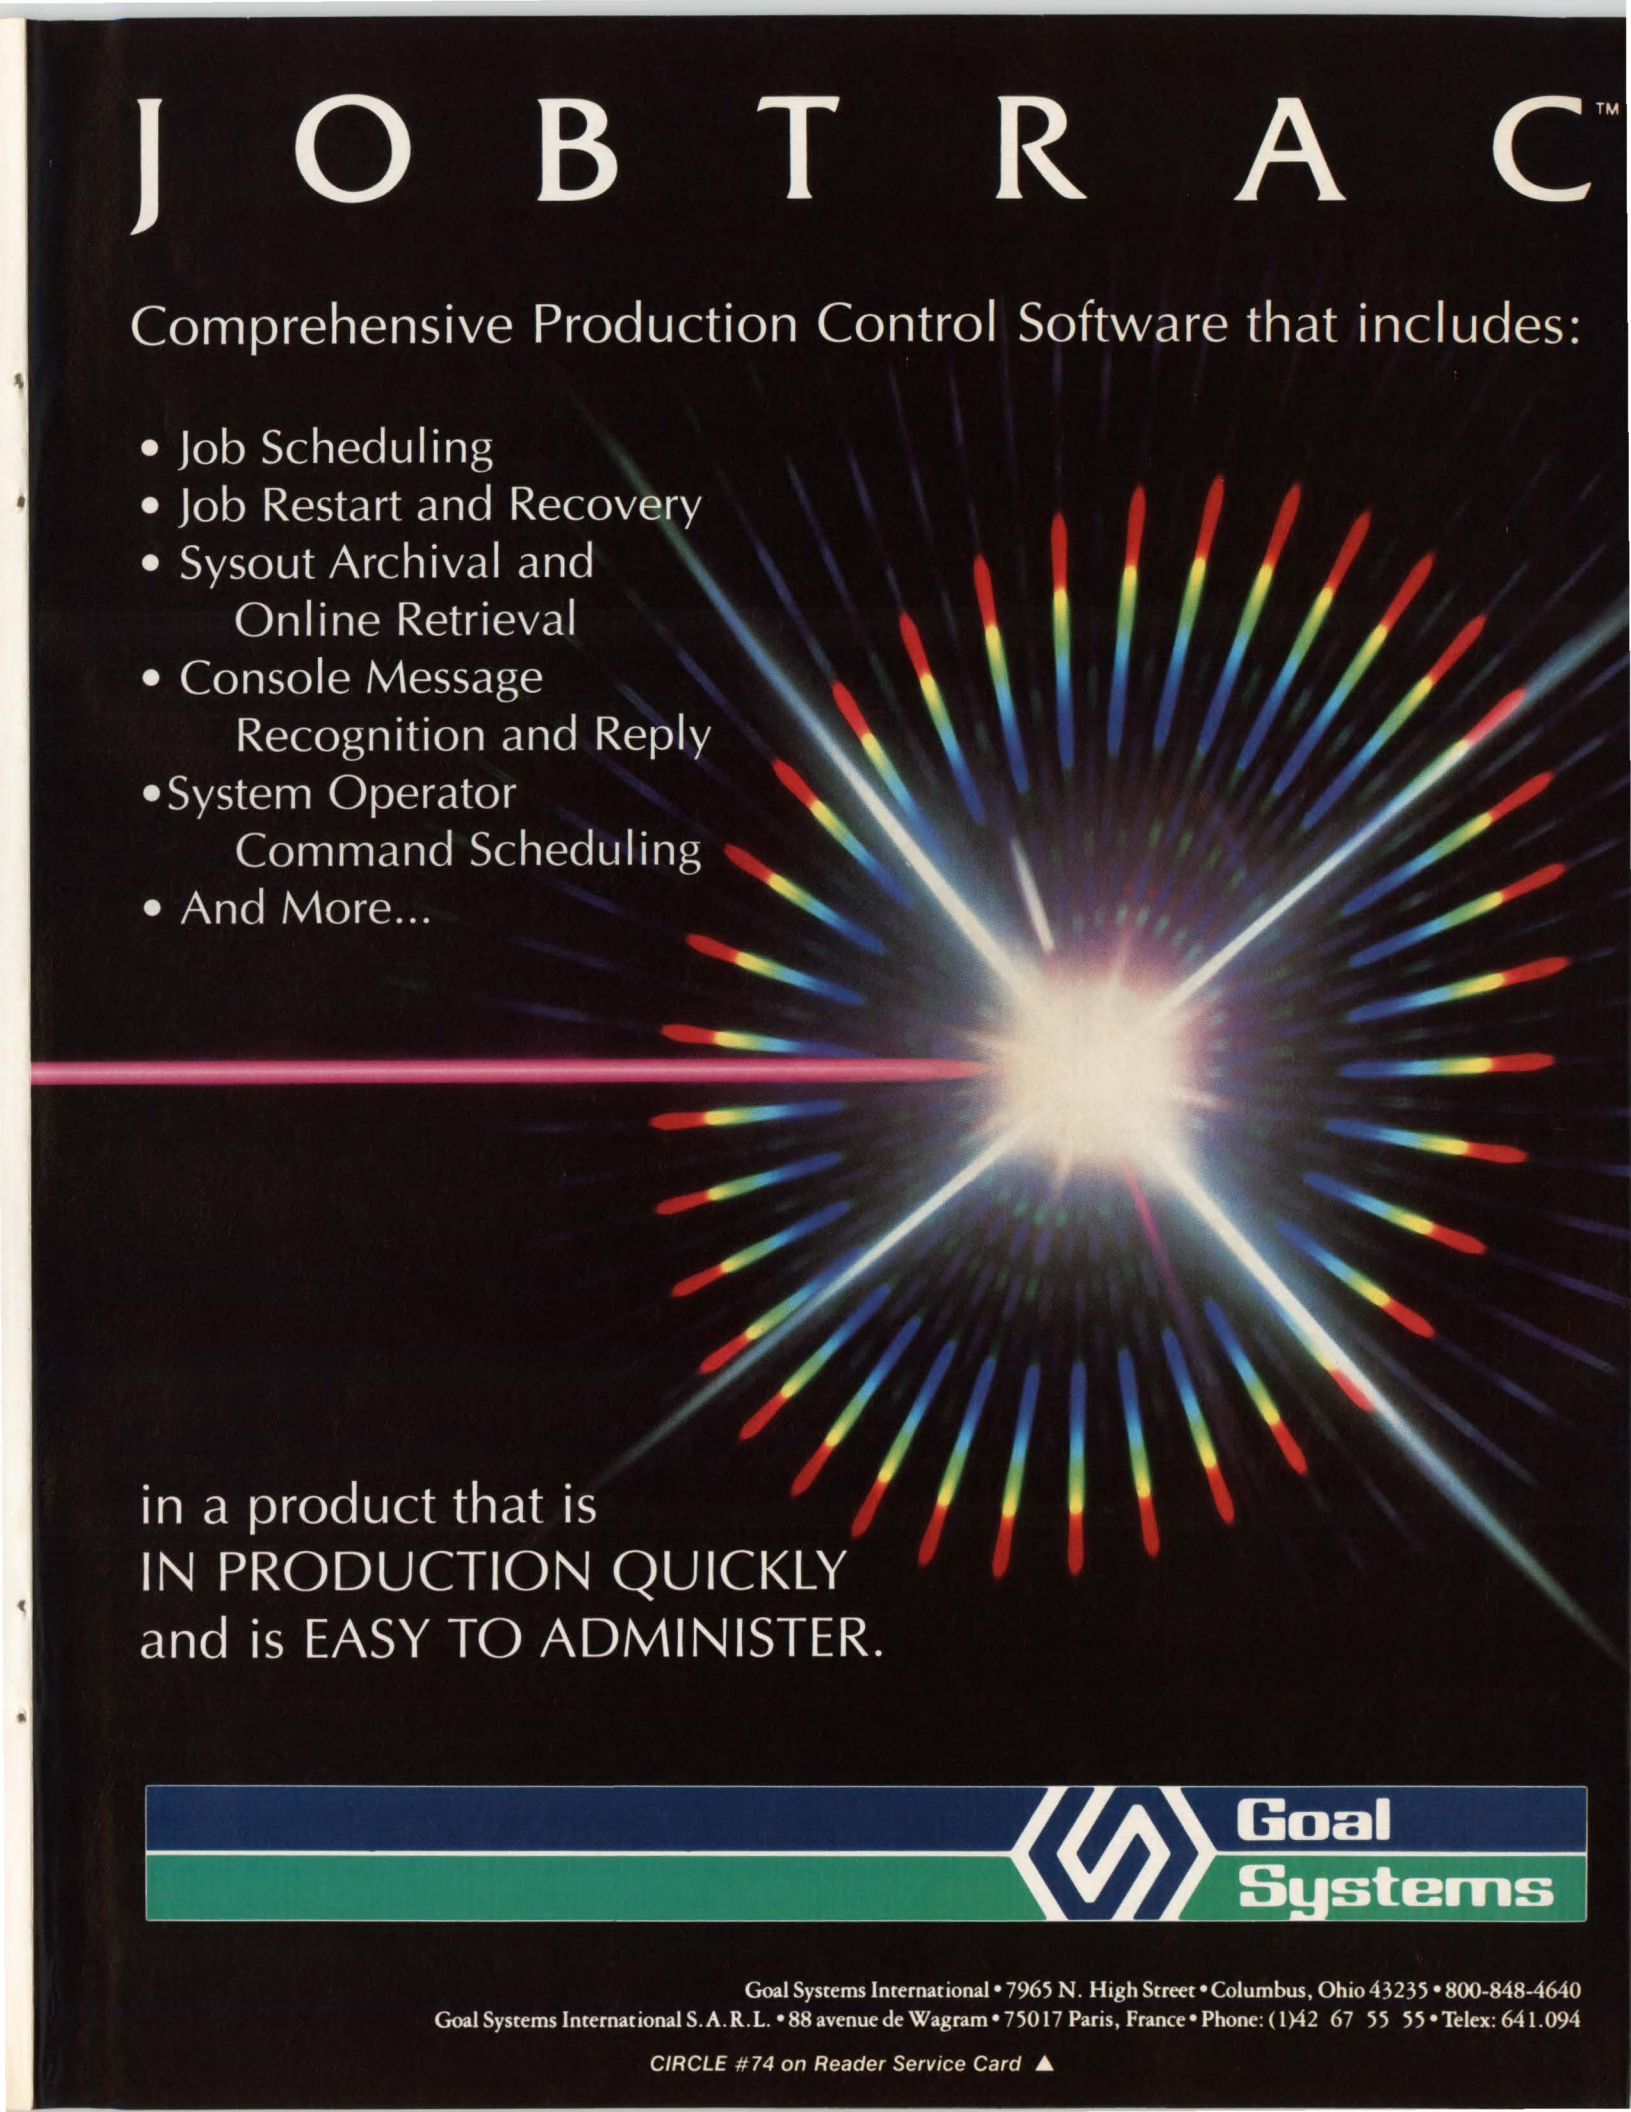

in a product that is
IN PRODUCTION QUICKLY
and is EASY TO ADMINISTER.

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 #74 on Reader Service Card ▲

CICS Dump Processing With DFHPDX

By Phyllis Donofrio

Each type of CICS dump has a purpose and a place in the problem determination process. Unformatted dumps have been used infrequently by many systems programmers, mainly because there was no facility to take the unformatted contents and make any sense out of them. Some installations have written their own routines or TSO Command Lists (CLISTs) to process the dump, to locate CICS control blocks and contents. IBM has provided a facility called IPCS to process unformatted dumps, but the product cannot format any CICS resources. There is now a facility in CICS to format both SDUMPs and SYSMDUMPs from CICS Release 1.7 and above. This article will explain how to use the facility and how to interpret the output. The output can be extremely valuable and save hours in analyzing problems.

CICSDATA

The command or verb in AMDPRDMP that requests CICS formatted output from an unformatted dump is CICSDATA. Before using this command in AMDPRDMP, two processes need to be completed. The first is to install the program or exit into a CICS loadlib for execution. The second is to make this command name known to AMDPRDMP.

The steps for each of these two processes are different, depending on the release of CICS and the release of MVS that was installed.

Installing DFHPDX Into CICS/MVS 2.1

The instructions to install this facility can be found in the *CICS/MVS Operations Guide* for CICS/MVS 2.1. IBM incorporated this facility into CICS/MVS 2.1, so after installation of the product is complete, module DFHPDX will reside in the CICS loadlib. The only additional steps necessary are to define the verb, CICSDATA, to AMDPRDMP so that when executing the command it will correlate the program name to that command.

Installing DFHPDX Into CICS 1.7

IBM released the PRDMP facility after CICS 1.7 had been available for almost two years. In order to incorporate the function into the product, therefore, IBM shipped the facility as a PTF. The APAR PL18949 or maintenance from 8803 PUT level will apply the function into a CICS 1.7 system. All documentation to install the function and utilize the product in a CICS 1.7 system is within the PTF cover letter. Since CICS manuals are not updated in CICS 1.7, any documentation change is made available via the service process (SMP apply of a PTF). The PTF adds the load module DFHPDX into the CICS load library for execution.

Installing CICSDATA Into MVS/XA 2.2

This step applies to support for the PRDMP facility in MVS/XA 2.2. AMDPRDMP has a table called the Exit Control Table. This table must be modified to contain the CICS exit that will be invoked to format the dump. In MVS/XA 2.2, a new member was created in SYS1.PARMLIB. This member, BLSCECT, contains an entry for all AMDPRDMP control statements and IPCS verb exits. The member is created by the installation of MVS and contains all verbs shipped and supported by the base product. It can be modified to add additional functions. Figure 1 contains a portion of this member. Notice that an item has been added to the end of the EXIT entries for the program DFHPDX (1). This

entry defines the Exit Program (EP) DFHPDX to be invoked by the VERB CICSDATA. This is a simple and clean way to add the function for the exit. In older releases of MVS, the process becomes much more complicated. After installing DFHPDX into the CICS loadlib and updating the BLSCECT member in SYS1.PARMLIB, you are ready to invoke the program.

FIGURE 1

SYS1.PARMLIB Member BLSCECT With CICSDATA Verb

```
   EXIT EP(HASPBLKS) VERB(JES2)     /* JES2 analysis                    */ 01850000
   EXIT EP(IATABPR)  VERB(JES3)     /* JES3 analysis                 ∂P3C*/ 01900000
   EXIT EP(IEAVTREF) VERB(LOGDATA)  /* LOGREC formatter                 */ 01950000
   EXIT EP(IEEMB817) VERB(MTRACE)   /* Master TRACE formatter           */ 02000000
   EXIT EP(IEAVNUCM) VERB(NUCMAP)   /* Nucleus mapping routine          */ 02050000
   EXIT EP(ISGDPDMP) VERB(Q)        /* Alias for QCBTRACE               */ 02100000
   EXIT EP(ISGDPDMP) VERB(QCBTRACE)  /* GRS ENQ formatter               */ 02150000
   EXIT EP(IARRDMP)  VERB(RSMDATA)  /* RSM analysis                     */ 02200000
   EXIT EP(IRARMFMT) VERB(SRMDATA)  /* SRM analysis                     */ 02250000
   EXIT EP(AMDSAFCM) VERB(SADMPMSG)  /* SADMP console message dump      */ 02300000
   EXIT EP(IEAVTFSD) VERB(SUMDUMP)  /* Summary dump formatter           */ 02350000
   EXIT EP(BLSQSUM1) VERB(SUMMARY)  /* Summary processor                */ 02400000
   EXIT EP(ASRSYMV)  VERB(SYMPTOM)  /* SYMREC symptom formatter         */ 02450000
   EXIT EP(ASRSYMV)  VERB(SYMPTOMS)  /* SYMREC symptom formatter        */ 02500000
   EXIT EP(IEDPRDMP) VERB(TCAMMAP) AMASK(X'00FFFFFF')  /* TCAM          */ 02550000
   EXIT EP(IEAVETFC) VERB(TRACE)    /* System TRACE formatter           */ 02600000
   EXIT EP(IKJVETSO) VERB(TSODATA)  /* TSO analysis                  ∂D1A*/ 02637500
   EXIT EP(IGVSFMAN) VERB(VSMDATA)  /* VSM analysis                     */ 02675000
   EXIT EP(ISTRAFD1) VERB(VTAMMAP) AMASK(X'00FFFFFF')  /* VTAM          */ 02712500
①  EXIT EP(DFHPDX) VERB(CICSDATA) /* CICS                             */ 02712600
/* ---------------------------------------------------------------------- */ 02750000
/* TCB formatting exits--invoked in the order listed                      */ 02800000
/* ---------------------------------------------------------------------- */ 02850000
                                                                          02900000
   EXIT EP(IECDAFMT) TCB            /* Data Management TCB exit         */ 02950000
   EXIT EP(IECIOFMT) TCB            /* IOS TCB exit                     */ 03000000
   EXIT EP(IEAVTFMT) TCB            /* RTM TCB exit                     */ 03050000
   EXIT EP(IEAVSSA1) TCB            /* Vector feature TCB exit          */ 03075000
/* ---------------------------------------------------------------------- */ 03100000
/* CBSTAT exits--invoked in the order listed                              */ 03150000
/* ---------------------------------------------------------------------- */ 03200000
                                                                          03250000
   EXIT EP(IEAVTRCA) CBSTAT(TCB)    /* RTM TCB status exit              */ 03300000
   EXIT EP(IEAVG701) CBSTAT(TCB)    /* COMMTASK TCB exit for WTORs      */ 03350000
   EXIT EP(IEAVTRCA) CBSTAT(ASCB)   /* RTM ASCB status exit             */ 03400000
   EXIT EP(IRARMCBS) CBSTAT(ASCB)   /* SRM ASCB status exit             */ 03450000
/* ---------------------------------------------------------------------- */ 03500000
/* ANALYZE exits--invoked in the order listed                             */ 03550000
/* ---------------------------------------------------------------------- */ 03600000
                                                                          03650000
   EXIT EP(IEAVESLX) ANALYZE        /* Supervisor lock analysis         */ 03700000
   EXIT EP(IOSVFMTH) ANALYZE        /* IOS I/O contention analysis      */ 03750000
   EXIT EP(ISGDCONT) ANALYZE        /* GRS ENQ contention analysis      */ 03800000
```

FIGURE 2

Sample Job To Dump Contents Of AMDPRECT

```
//DUMPTJ JOB (,,9),SYS.PROG,MSGCLASS=X,CLASS=N,                          
// MSGLEVEL=(1,1)                                                        
//DUMPT   EXEC PGM=AMASPZAP                                   00000040
//SYSPRINT DD SYSOUT=*                                        00000050
//SYSLIB DD DSN=SYS1.LINKLIB,DISP=SHR                                   
//SYSIN DD *                                                  00000080
DUMPT AMDPRECT                                                00000081
```

Installing CICSDATA Into MVS Releases Prior To MVS/XA 2.2

In MVS Releases 2.1.x or prior to MVS 2.2, the Verb Exit Table is in a SYS1.LINKLIB module called AMDPRECT. Since this module was created by MVS and no SYS1.PARMLIB member is available, you must modify the module with the proper information via a facility in MVS called SUPERZAP. This facility allows you to overlay the existing contents of a module with new values and is documented in *MVS System Programming Library/Service Aids*. The steps necessary would be as follows:

First, locate an empty slot in the table by running the DUMPT function of AMASPZAP. This step dumps, in hexadecimal format, the contents of the module that needs to be modified. Figure 2 shows a sample job to locate an empty slot in the table. Output similar to that in Figure 3 will be produced. This is a portion of the output of the DUMPT request. It displays the contents of AMDPRECT from which you will need to locate an empty slot. Scanning through the output, you will see the first empty slot at location 0258 (1). These five fullwords contain the sequence necessary to reuse for the CICS exit. Insert both the command or verb to use plus the exit name for AMDPRDMP to execute. Remember how easy it was in MVS/XA 2.2, merely updating a parameter library? Well, instead, it is necessary to overlay the module with the hexadecimal representation of the same information. Not quite as easy and definitely not clean. Figure 4 contains a sample job that would update the module with the information necessary. Identify the module name to ZAP, AMDPRECT (1). The

FIGURE 3

DUMPT Output Of AMDPRECT

```
000160  C4E2 D5E6 C4D4 D740 C4E7 D9D9 D3D4 F5F0  0080 0000 C9D9 D3D4 4040 4040 4040 4040  *DSNWDMP DXRRLM50*
            CLC       XC        MVCK MVZ                      MVZ  STH  STH  STH  STH   *....IRLM        *
000180  4040 4040 0000 0000 4040 4040 4040 4040  C1C4 E8C8 C4C6 D4E3 0400 0000 C4C1 C5C4  *    ....        *
        STH  STH            STH  STH  STH  STH                 NC   SPM                  *ADYHDFMT....DAED*
0001A0  C1E3 C140 C9C5 C5D4 C2F8 F1F7 0000 0000  D4E3 D9C1 C3C5 4040 C2D3 E2D8 E2E4 D4F1  *ATA IEEMB817....*
                       MVO                       NC   MVCK      STH                 NC    *MTRACE  BLSQSUM1*
0001C0  0000 0000 E2E4 D4D4 C1D9 E840 C1D4 C4E2  C1C6 C3D4 0000 0000 E2C1 C4D4 D7D4 E2C7  *....SUMMARY AMDS*
                  NC                                                 XC                  *AFCM....SADMPMSG*
0001E0  C9C5 C4D7 D9C4 D4D7 0080 0000 E3C3 C1D4  D4C1 D740 C9C1 E3C1 C2D7 D940 0080 0000  *IEDPRDMP....TCAM*
                  MVCK NC                        NC   XC                 MVCK             *MAP IATABPR ....*
000200  D1C5 E2F3 4040 4040 C9C5 C1E5 E2E2 C1F1  8000 0000 4040 4040 4040 4040 C8C1 E2D7  *JES3     IEAVSSA1*
        MVN       STH  STH                       SSM       STH  STH  STH  STH            *....         HASP*
000220  C2D3 D2E2 0000 0000 D1C5 E2F2 4040 4040  C1E5 C6D9 C4C6 D4E3 0000 0000 C1E5 D4C4  *BLKS....JES2     *
            MVC            MVN       STH  STH                 NC                  NC     *AVFRDFMT....AVMD*
000240  C1E3 C140 C9D2 D1E5 C5E3 E2D6 0000 0000  E3E2 D6C4 C1E3 C140 4040 4040 4040 4040  *ATA IKJVETSO....*
                  MVN                            OC                  STH  STH  STH  STH  *TSODATA         *
000260  0000 0000 4040 4040 4040 4040 4040 4040  4040 4040 0000 0000 4040 4040 4040 4040①  *....            *
                  STH  STH  STH  STH  STH  STH   STH  STH                 STH  STH  STH  STH  *    ....        *
000280  4040 4040 4040 4040 0000 0000 4040 4040  4040 4040 0000 0000                        *        ....    *
        STH  STH  STH  STH                 STH  STH  STH  STH                              *    ....        *

**CCHHR- 0056000A06   RECORD LENGTH- 000000          MEMBER NAME  AMDPRECT

AMA113I COMPLETED DUMP REQUIREMENTS

AMA100I AMASPZAP PROCESSING COMPLETED
```

FIGURE 4

Sample Job To ZAP AMDPRECT With CICSDATA Verb

```
//ZAPJOB JOB (,,9),SYS.PROG,MSGCLASS=X,CLASS=N,
// MSGLEVEL=(1,1)
//ZAPCICS EXEC PGM=AMASPZAP                                    00000040
//SYSPRINT DD SYSOUT=*                                         00000050
//SYSLIB DD DSN=SYS1.LINKLIB,DISP=SHR
//SYSIN DD *                                                   00000080
NAME AMDPRECT                                                  00000081
VER  0258 40404040,40404040,00000000,40404040,40404040         00000085
REP  0258 C4C6C8D7,C4E74040,00000000,C3C9C3E2,C4C1E3C1         00000087
```

next line is a VER or verification statement to the program. Since it will be necessary to modify MVS operating system code, verification is first performed for the correct location.

The VER statement verifies that at offset 0258 into this module, the following contents already exist. If the VER statement does not match the actual contents within the module, the ZAP fails and the modification does not take place. Assuming that the VER matches, the next line contains the REP or replacement entries. In this example, we are replacing at location 0258 the hexadecimal entries contained in the five fullwords. The first two fullwords, C4C6C8D7,C4E74040, will insert the program name DFHPDX. The last two fullwords, C3C9C3E2, C4C1E3C1, will insert the command CICSDATA. As you can see, moving to MVS/XA 2.2 will provide a much easier facility to add entries into this table.

Executing DFHPDX Against CICS Unformatted Dumps

After following the appropriate instructions to install the program into a CICS load library and making the verb known to the AMDPRDMP facility in MVS, using this product is as easy as running a batch job with the appropriate JCL. Of course, AMDPRDMP is fully documented in the IBM manual, *Service Aids*. However, since the MVS manual knows nothing about the additional verb, it will not contain instructions on how to execute the instructions. Figure 5 contains the CICSDATA command and the available parameters. The parameters merely request additional control blocks in the CICS dump. Utilize the parameter for the appropriate CICS data area and it will appear in the output listing. If the command CICSDATA is used with no options, all available control blocks are reported. If analyzing a problem dealing with terminal failures, you may wish to request only the data areas relevant to that problem. Depending on the size of the region, and therefore the amount of storage that was dumped, the output from this job can be significant. Choose the output that will be most relevant to your problem. If you are diagnosing a difficult or unknown problem, you will probably decide to request everything. The amount of lines produced are still insignificant compared to printing the complete unformatted dump.

Figure 6 contains a sample job to invoke this facility. Of course, the JCL will vary based on the installations naming conventions. The most critical items in this job are the STEPLIB dataset containing the CICS module and the dataset pointing to the unformatted dump (1). This dataset must be the SYS1.DUMP dataset or the SYSMDUMP dataset containing the unformatted dump. After executing this job, the output listing will contain the CICS data areas requested from the options. Following are several examples of the output created by this facility.

Figure 7 contains the first page of output from DFHPDX. Prior to this page in the listing will be several pages from AMDPRDMP; however, most of the contents pertain to the result of processing the unformatted dump and identity of the CPU hardware involved with the dump. Figure 7 contains the first page that would be relevant to CICS contents. Line one on every page will contain the following information.

- Title of the dump either keyed by the operator when requesting the dump or the title from the SYS1.DUMP dataset.
- Module in control at the time of the failure. The module IEAVTSDT is the MVS module invoked by dump control if the operator requested the dump from the console.
- The time and date stamp logged at the time of the dump.

The first CICS control blocks listed are the CSA (4) and the CSA Optional Features List (5). These and other areas within the listing are identical to CICS data areas that could be requested by a formatted dump. Remember, however, that although the output is formatted, a complete dump of the entire address space has been recorded. The areas displayed on this report are only a portion of the storage that is available within the unformatted dump.

Invocation of the CICSDATA verb will, by default, produce formatted output of the CSA, the CSA Optional Features List, the SIT, Abend Trace Table, Trace Table, Registers and Static Storage. Using additional keywords will produce formatted control blocks corresponding to the keyword. For example, a request for KCP will also produce a task summary at the time of the dump and all active and suspend chains with corresponding user storage. The advantage of this facility, as opposed to a normal CICS formatted dump, is that requests can be made for portions of the dump that are relevant to the problem. Even a formatted dump of a normal CICS region can be a substantial amount of paper. This PRDMP facility provides selective format and review of control blocks, producing a much smaller report.

One of the most important pages within this formatted listing can be found far

FIGURE 5

PRDMP Control Keywords

JOB=jobname/CURRENT

PCT	PROGRAM CONTROL TABLE
PPT	PROCESSING PROGRAM TABLE
SCP	STORAGE CONTROL PROGRAM CHAINS
MCT	MONITOR CONTROL TABLE
DLI	DLI INTERFACE AND CONTROL BLOCKS
KCP	TASK SUMMARY, DISPATCHER QUEUES AND ENQ BLOCKS
TSP	TEMPORARY STORAGE QUEUES
ICE	INTERVAL CONTROL ELEMENTS
FCT	VSAM SUBTASK AND FILE CONTROL TABLE
TCT	TERMINAL CONTROL TABLE AND ENTRIES
MRO	MULTIPLE REGION OPERATION CONTROL BLOCKS
TDP	TRANSIENT DATA CONTROL BLOCKS
XRF	XRF CONTROL BLOCKS

FIGURE 6

Sample Job To Execute CICS Formatted Dump Program

```
//CICSDATA JOB (,,,9999),SYS-PROG,CLASS=N,MSGCLASS=X
//*                                                            00000002
//* FOR PRINTING SYSMDUMPS AND SYS1.DUMPXX WITH CICS EXIT       00000003
//*                                                            00000004
//PRINT    EXEC PGM=IKJEFT01,PARM=AMDPRDMP
//STEPLIB DD DSN=CICS210.LOAD,DISP=SHR
//SYSPRINT DD SYSOUT=*
//INDEX    DD SYSOUT=*                                         00000008
//SYSUT1   DD SPACE=(CYL,(50,10)),UNIT=SYSDA
//SYSTSPRT DD DUMMY                                            00000010
//SYSTSIN  DD DUMMY                                            00000011
//SYSIN    DD *                                                00000012
  CICSDATA KCP
  END                                                          00000170
//TAPE     DD DSN=SYS1.DUMP00,DISP=SHR                  (1)
//                                                             00000190
```

Response Time is Money.

Poor CICS response time is expensive. The longer users wait, the less they get done, the more your bottom line suffers, and the more you get blamed. But good response time can also be expensive - when it's purchased through more hardware or overworked systems staff.

No guesswork
Candle helps keep CICS response time and your budget at a minimum with software that precisely pinpoints the causes of poor CICS performance. Our end-to-end response time feature even tells you whether the problem is in the host or in the network. And that improves the response time of your staff.

Faster solutions
The OMEGAMON® family of products for CICS keeps your user service levels on target by detecting availability threats and slowdowns immediately, analyzing the cause for you, and recommending the solution. A single Status Monitor™ screen keeps you on top of your entire CICS network - across CPUs and across geographic locations. OMEGAMON eliminates the time-consuming process of analyzing irrelevant resource data by doing the analysis for you.

Our products also help you budget. They give you trending and capacity planning information so you can forecast future needs and make maximum use of existing resources. With software, not by throwing more iron at your problems before it's really needed.

Total support
As for our response time, Candle's support team is available round-the-clock, around the world, every day of the year to answer your questions and help you get the best performance from your CICS system. And Candle Education helps you improve your own performance.

Is your response time costing you money? It's worth your time to find out by calling your Candle Account Representative or Terry Forbes today at **(800) 843-3970.**

!Candle®

CIRCLE #120 on Reader Service Card ▲

FIGURE 7

Formatted Dump Output: CSA And CSAOPFL

```
 (1)                                                                    (2)                        (3)
CICS TEST                                                    MODULE IEAVTSDT DATE 11/03/88  TIME 08.49.24  PAGE 00000004

-- FORMATTING CONTROL BLOCKS FOR JOB CICS21TS

 REGISTERS 0-15 00A93050

   0000  00000001 FFFF58C4 005D4230 005D2EE0  005CD948 005CE948 0052EDE8 00000001  *.......D.)...)...*R..*Z....Y....*   00A93050
   0020  0000A73C 0053D348 00000000 005877D8  005D4130 005D2670 00000000 80ABA130  *..X...L........Q.)...).........*   00A93070

 DFHCSA 005D2670 (4)

   0000  00000000 00005F98 00031098 705CBA12  005D2FC0 00000000 FFFFFFFF 005D4230  *......¬Q...Q.*...)...........)..*   005D2670
   0020  FFFFFFFF 405CB9DA 005CC9DA 00000000  00000000 000B4000 0052EDE8 FFFFFFFF  *.... *...*I...........  ....Y....*   005D2690
   0040  005877D8 00172008 0010020C 005D3150  0849239F 00040840 07D00100 00000000  *...Q.........).  .......  .......*   005D26B0
   0060  003077CC 00012C00 00ABA130 00006068  0000A000 00006000 005E1000 0988308F  *..............-.......-..;...H..*   005D26D0
   0080  005D4130 E8FFFFFE 0000001E 005D5888  000000D0 00000000 0003F030 E738E721  *.)..Y........).H...........0.X.X.*   005D26F0
   00A0  00000000 000402A0 000400E0 00040230  00040380 005D2EE0 005D3150 C500FF00  *.......................)...). E...*   005D2710
   00C0  000A0003 00000468 005D2AD8 00C800B0  00000000 00000000 00ABA130 00000000  *.........).Q.H..................*   005D2730
   00E0  405CB9D4 00584FF8 0057F4B0 005C1BA8  0054B38C 00587DA8 0057779C 005606F8  * *.M..|8..4..*.Y......'Y.......8*   005D2750
   0100  0056441C 00000000 0053D6AC 00531654  00000000 005C732C 015BABF0 00534960  *..........0..........*...$.0...-*   005D2770
   0120  00000000 00000000 005877D8 00569668  00560450 00446378 00063FF0 005D4C28  *...........Q..0...  ......0.)<.*   005D2790
   0140  00000000 00531680 00000000 00000000  00000000 00530EA0 005D5838 00000000  *..........................).....*   005D27B0
   0160  00000000 00000000 00000000 5CCFBE60  0070FA1D 02000100 00000000 00000000  *............*..-.................*   005D27D0
   0180  00000000 00000000 00000000 000071F0  00000000 00000000 00000000 005D2608  *...............0..............)..*   005D27F0
   01A0  07FE58F0 D19C07FF 00916766 00916892  00177000 000C0000 E6D6D9D2 C1D9C5C1  *...0J....J...J.K........WORKAREA*   005D2810
   01C0  00000000 001C0000 003C004C 00189C00  00492900 003B9700 00000000 00000000  *...........<..........P.........*   005D2830
   01E0  00000000 000C0000 0C000C00 0C00000C  00000C00 000C0000 0C00000C 00000C00  *................................*   005D2850
   0200  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *................................*   005D2870
   0220  -     043F LINES SAME AS ABOVE                                                                                  005D2890
   0440  00000000 00000000 00000000 00000040  C100000C F1F16?F0 F361F8F8 001B6120  *............... A...11/03/88../.*   005D2AB0
   0460  0052D000 00000000                                                         *........                        *   005D2AD0

 CSAOPFL 005D2AD8 (5)

   0000  00000000 00000000 00530130 005505A8  005593A2 00000000 00582128 005821BC  *...............Y..LS............*   005D2AD8
   0020  00549AEC 00549498 00000000 00000000  00000000 00581A58 00000000 00000000  *......MQ........................*   005D2AF8
   0040  00000000 00000000 00000000 00000000  005409C4 00000000 00000000 005D4D48  *...................D.........)(.*   005D2B18
   0060  805D4D84 01F4F15B 00000000 0000FE00  00550138 0088308F 005774D8 005D2670  *.)(D.41$.............H.....Q.)..*   005D2B38
   0080  005D4E28 00584B28 0054D1A8 8220ACC0  00000000 00000000 00000000 00000000  *.)+.......JYB...................*   005D2B58
   00A0  00000000 00000000 005CF4CC 005D1798  005C0438 0057C6B8 00564158 00569018  *.........*4..).Q.*....F.........*   005D2B78
   00C0  00586E88 0057D2E8 005C2BE8 005CEEB8  005843D8 005C8328 0055BA48 0053D168  *..>H..KY.*.Y.*.....Q.*C.......J.*   005D2B98
   00E0  0054D0F8 005501E8 00551CA8 0055F768  00000000 00000000 005D2D80 005C4398  *...8...Y...Y..7..........)...*.Q*   005D2BB8
   0100  005B11F8 00584958 00584648 005B94A8  005B2C98 00000000 005BA66C 00000000  *.$.8.........$MY.$.Q.....$W%....*   005D2BD8
   0120  005AA8C8 005A39C8 005A05D8 00553E82  00554C72 0055DB22 0055C6A8 005585D2  *..YH...H...Q...B..<.......FY..EK*   005D2BF8
   0140  0055EF62 0055A162 00556D02 00553738  00557352 0054F478 0057E3FC 0000779C  *.........._..........4...T.....*   005D2C18
   0160  00536078 0059881C 00020BA8 005C330C  005C8508 00000000 00000000 005C9F48  *..-...H....Y.*...*E..........*..*   005D2C38
   0180  00538C08 005808A8 00000000 005C4C9C  00000000 00551D7C 00547E8C 0054690C  *.......Y.....*<.........a..=.....*   005D2C58
   01A0  0059D9A8 00000000 002CD008 00000000  00000000 0059CE48 00530E00 005D2E00  *..RY..........................)..*   005D2C78
   01C0  00534FB0 00078374 0000A000 0053D348  00000000 00531268 00000000 00040460  *..|...C.......L................-*   005D2C98
   01E0  00000000 00046188 000057B8 00000000  40000000 00000000 00000000 00000000  *....../H.......  ...............*   005D2CB8
   0200  00000000 00000000 80C00000 0000A000  00453E68 C4D7C1C4 C3C9E200 00000000  *....................DPADCIS.....*   005D2CD8
   0220  0AD80AD7 00000000 00000000 00000000  00000000 00000000 00535DC0 00537000  *.Q.P.........................)....*   005D2CF8
   0240  9F5CCFB6 4C384400 9F041030 403AFC00  00000000 00000001 00006020 000A7420  *.*..<.......  ..............-.....*   005D2D18
   0260  00000001 80800000 00045008 40404040  8000DA98 00000000 00000000 00000000  *.........  .       ...Q..........*   005D2D38
   0280  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *................................*   005D2D58
   02A0  00000000 00000000 0059CFB8 00000000  00000000 00000000 00000000 00000000  *................................*   005D2D78
   02C0  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *................................*   005D2D98
   02E0  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *................................*   005D2DB8
```

within the report.

Figure 8 contains a page that will be of incredible importance when diagnosing difficult failures. One of the critical pieces of information required in problem determination is the identity of the active and suspend chains. This information contains critical information pertaining to the failure. This example shows the Summary of Tasks on the Active Chain (1), as well as the suspend chain directly following (2).

Of course, IPCS can be used to locate pointers and chase these chains. Many systems programmers have done this themselves for years because no facility was available. This could, and has, taken hours to locate in a heavily loaded CICS region with many transactions active. This job can now be processed against an unformatted dump to provide a visual look at exactly what was happening in the region at the time of the failure.

The first summary identifies the active Dispatch Control Area (DCA) chain and some associated items (1). The items listed for you are:

- The address of this DCA entry within the dump (3)
- The task number of the task, if available (4)
- The address of the TCA user area for this task (5)
- The address of the event control area for this task (6)
- The contents of the event control area status (7)
- The contents of the task control dispatch indicator (8)
- An interpretation of the above indicator (9).

From these entries, tasks on both the active and suspend chain are available with a small bit of information about each. This may appear insignificant but can provide a wealth of information about the status of the region and currently executing tasks at the time of the failure.

Following these summaries is a complete listing of each DCA (10), the corresponding TCA User Area (11) and TCA System Area (12). The next several pages in the report list the complete chains both active and suspend. From these areas, it is possible to inspect contents of task control indicators to determine their potential relevance to the problem. This information, with minimal effort, can provide a quick snapshot of conditions and save enormous amounts of time in the problem determination process. Of course, if initial investigation proves to be inadequate, this information can supplement additional facilities. Since a full unformatted dump was recorded on DASD, IPCS can investigate additional areas of the dump. The combination of DFHPDX to get a quick look at the problem and IPCS to investigate any other area of storage necessary, provide two powerful tools to debug many different types of problems.

Invocation Of VERBEXIT Within IPCS

While the program DFHPDX was orig-

IT WOULD TAKE THE AVERAGE PROGRAMMER 2 HOURS TO FIND THE ERRORS ON THIS SCREEN.

VERIFY® WOULD BE FINISHED BY NOW.

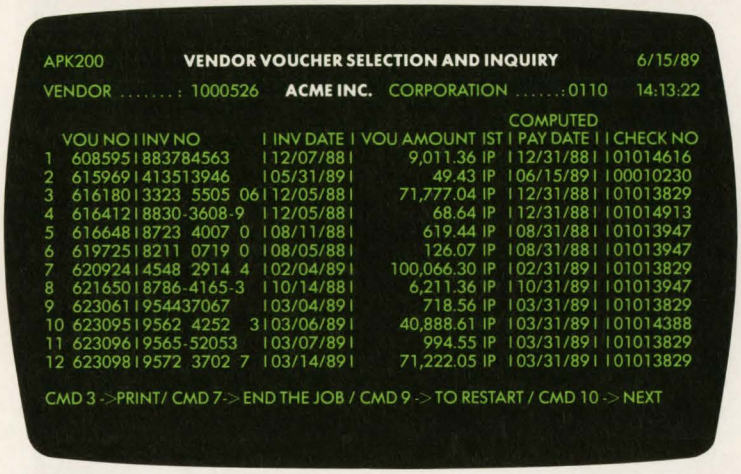

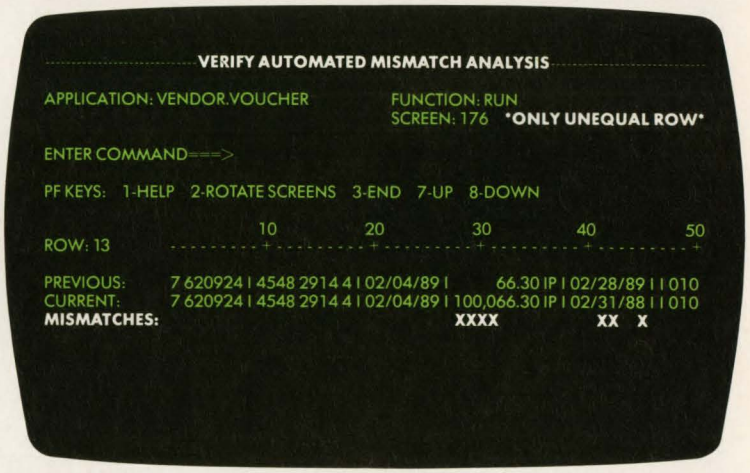

In fact, in the 84 seconds it will take you to finish reading this ad, VERIFY could log a test script using the current version of a program, rerun it with an updated version of the program, and then automatically identify and resolve all errors and unexpected results. VERIFY prevents production system surprises. Immediately. Interactively.

And it performs just this efficiently on a variety of CICS tests including unit, regression, stress, integration, concurrency, and migration—all of which can be performed using either an on-line or batch approach.

It makes the job easier, more accurate, faster, and maybe even a little more fun. So the average programmer would not only be finished by now, he or she would be testing something else—so your systems would need less testing in the future. Sort of a virtuous cycle.

If you're in a hurry to do nothing, call or write us at Two Executive Drive, Fort Lee, NJ 07024.

800-642-0177

In Canada, call 201-592-0009.

On·Line Software
INTERNATIONAL

On·Line Software. The Safe Buy.

All our products are offered with a lifetime trade-in guarantee so that the money you spend today is always available to meet your changing needs tomorrow.

On·Line Software offers consulting, education and software products. We specialize in CICS, DB2, and CASE technologies for application development.

MJYTG9

CIRCLE #196 on Reader Service Card ▲

FIGURE 8

Formatted Dump Output: Task Summary And Chains

```
CICS TEST                                          MODULE IEAVTSDT DATE 11/03/88  TIME 08.49.24  PAGE 00000030

== SUMMARY OF TASKS ON THE ACTIVE CHAIN (1)

DCAADDR  TASK ID/ADDR   ECAADDR  ECA FLAGS      (9)

040380  00023/0A8190  530970  00  8040   WAIT ON SYSTEM ECB
040310  00188/0A5190  0B4240  00  8840   WAIT ON CICS RELATED ECB IN DFHFCP   AT   00290C
0403F0  00036/04F190          --  4040   WAIT ON LIST OF ECBS
040070  00016/04B190  0461EC  00  8840   WAIT ON CICS RELATED ECB IN DFHVAP   AT + 000152
0401C0     JJJ/08C190  5D5838  00  8040   WAIT ON SYSTEM ECB
5D2EE0     TCP/5D3150          --  4040   **** CURRENT TASK ****
(3)        (4)         (5)     (6) (7)    (8)

== SUMMARY OF TASKS ON THE SUSPEND CHAIN (2)

DCAADDR TASK ID/ADDR  ECAADDR ECA FLAGS

0400E0     JJJ/047190          1004   NON DISPATCHABLE
040150     JJJ/04A190          1004   NON DISPATCHABLE
040000  00174/04D190           1340   TERMINAL WAIT
040230     JJJ/051190          1004   NON DISPATCHABLE

== DCA-TCA ACTIVE CHAIN

DCA 00040380 (10)

    0000  81000070 00000000 00000000 00040310  005D2714 000A8190 80400080 00530970  *A.............)....A.. .....*  00040380
    0020  00000000 00000000 00000000 00000000  00000000 22008000 00000000 00000000  *................................*  000403A0
    0040  00000000 00000000 00000000 00000000  00000000 00000000 000A8524 00000001  *.........................E.....*  000403C0
    0060  00000000 00000000                                                         *........                      *  000403E0

USERTCA 000A8190 (11)

    0000  000A8000 00000000 00530954 005D2AD8  000A85E8 00530970 80400000 00818000  *.............).Q..EY..... ...A..*  000A8190
    0020  50553216 00000000 00000001 000A8710  000A85E8 40553010 000A8580 00000000  * ..............G...EY .....E.....*  000A81B0
    0040  000A8000 480001E8 00000000 00000000  00530954 00530970 50553216 00083030  *.......Y.............. ......*  000A81D0
    0060  40553428 005D4E28 00000000 000A8710  000A85E8 40553010 000A8580 000001C8  * ....)+........G...EY .....E....H*  000A81F0
    0080  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *................................*  000A8210
    00A0  -      00DF LINES SAME AS ABOVE                                          *                              *  000A8230
    00E0  00000000 00000000 24000000 00000000  00000000 00000000 00000000 00000000  *................................*  000A8270

SYSTCA 000A8000 (12)

    0000  8A040978 000A8000 000A8980 000A8980  8000023C 00530BBC 00040380 00000000  *.........I...I.................*  000A8000
    0020  0004D190 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *..J............................*  000A8020
    0040  00000000 000A8290 00000000 00000000  00000000 00000000 00000000 00000000  *......B.......................*  000A8040
    0060  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *................................*  000A8060
    0080  FE0A87D0 00000000 FE0A8570 FE0A8978  005D4E28 000A8580 00000000 00000000  *..G .....E...I..)+...E.........*  000A8080
    00A0  00000000 00000000 00000000 00000000  C3E2E2E8 00000000 00000000 00000000  *................CSSY..........*  000A80A0
    00C0  00000000 00000004 00000000 00000000  00000000 00000000 00000000 00000000  *................................*  000A80C0
    00E0  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *................................*  000A80E0
    0100  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *................................*  000A8100
    0120  C00010C 00530BBC 00000000 00040380  00000000 00000000 00000000 00000000  *..... .........................*  000A8120
    0140  00000000 00000000 000A82C0 0000025C  00000000 00000000 00000000 00000000  *..........B....*...............*  000A8140
    0160  00000000 00000000 00000000 00000000  00000000 00000000 00000000 00000000  *................................*  000A8160
    0180  00000000 00000000 00000000 00000000                                      *................               *  000A8180
```

inally intended for batch processing, the facility can be executed from within an IPCS session. This may be desirable, since all functions for debugging are then available in one place, IPCS. In addition, MVS/ESA will require that all dump analysis be done with IPCS. The syntax to invoke the VERBEXIT may change with MVS/ESA, but the IPCS facility will be the focal point.

When DFHPDX is called by the VERBEXIT command in IPCS, this program must be available to the TSO session. Most CICS customers do not place the CICS load library in their TSO procedures. The module, then, must be placed into a library that is defined to the TSO procedure or in a Linklist library. If the module cannot be located by the IPCS session, a message will be displayed indicating that the program cannot be found. The installation must place the CICS load library into the TSO procedure or copy the module into an existing library. IEBCOPY can be used to move DFHPDX into the desired LOADLIB.

Remember, however, that moving DFHPDX into another library will require additional maintenance considerations. After updating CICS with SMP/E maintenance, this module may be enhanced or changed by one or more PTFs. SMP/E will update the module in the CICS libraries but will not update that module in any other library. The new copy of DFHPDX will need to be moved into the alternate LOADLIB after applying maintenance. This step will need to be part of the SMP/E procedures when updating CICS.

After DFHPDX has been made available to the TSO session, the VERBEXIT command can be used within IPCS to execute the dump formatter. The syntax of this command is:

```
VERBEXIT CICSDATA 'parm,parm. . .
```

Execution of this facility with no parameters will utilize the defaults in the same way as a batch job. If a subset of the options is required, input the requested options into the parameter fields. In other words, if only a minimum of output is necessary, the following command can be used:

```
VERBEXIT CICSDATA 'KCP'
```

This will produce output of the CSA, Trace Table, task summary and all task storage. Additional options can be chosen for portions of the dump required.

The result of this command will produce the report output into the IPCS session. The report will appear identical to the printed output produced by a batch job. It is possible, however, to remain within the IPCS session and scan the output. For example, the results of the previous command would display formatted control blocks, which could be inspected. The FIND command could be used to locate specific entries within the report. The IPCS session can scroll forward or backward through the report to locate necessary items.

If the IPCS session is using the standard message routing defaults, NOPRINT TERMINAL, the output is routed to the terminal. When the IPCS session terminates, the output is lost. If the information needs to be retained, the IPCS defaults can be changed to route the output to the print dataset. IPCS allocates a file,

FETCH DESTROYS CICS STRESS

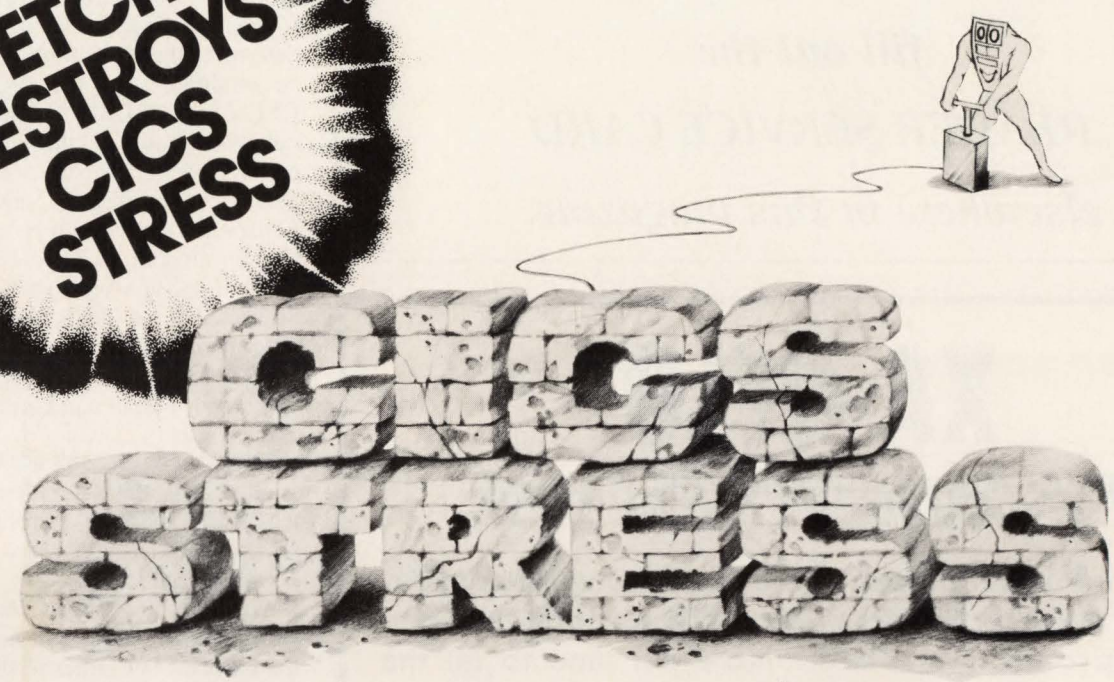

FETCH is a powerful solution to the problems that cause CICS stress. FETCH will dramatically improve CICS performance from day one. Just how good is it, really? Well, you could say FETCH is dynamite. Install FETCH on your CICS and get instant relief. Relief from CICS "bottleneck" problems like virtual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

IPCSPRNT, to be used during the session. Some installations specify this file in the IPCS CLIST during invocation. Check the CLIST that invokes IPCS to determine the allocation. The routing can be modified with the IPCS commands OPEN and CLOSE to allocate the print output explicitly. If the IPCSPRNT allocation identifies the standard SYSOUT = A specification, any output routed to the print dataset will be released at session termination to this output print class.

You can, however, specify this output be routed to an IPCS browse dataset. If the following DD is placed into the IPCS CLIST or specified in the IPCS OPEN command, output can be routed to a permanent file.

```
//IPCSPRNT DD DSN=IPCS.BROWSE,DISP=SHR
```

After this dataset is created, change the IPCS message definitions for the session with:

```
SETDEF PRINT NOTERMINAL
```

When the VERBEXIT command is used within IPCS, output will not be routed to the terminal but rather to the IPCSPRNT allocated destination. If this is a permanently-defined dataset, the output can be browsed with ISPF. If routed to a SYSOUT, it can be printed for later review.

The dump formatting utility will continue to be a valuable tool in problem determination and can be used in MVS/ESA within IPCS. The process is entirely supported in MVS/XA 2.2 and can be used with that release to gain experience with the function and syntax. Additional enhancements to both DFHPDX and IPCS will provide the tools necessary for debugging in future CICS offerings.

Since IPCS will continue to be enhanced, new functions will provide even more reason to utilize this facility. Utilization with the existing version of MVS/XA will give customers the foundation to take advantage of these enhancements when available.

ABOUT THE AUTHOR

Phyllis Donofrio is currently a senior systems programmer for special projects at Dun and Bradstreet Planned Services, 3501 Frontage Road, Tampa, FL 33607, (813) 289-1000 Ext. 2143. She adapted this article from her book, CICS: Debugging, Dump Reading and Problem Determination *(J. Ranade Series/McGraw-Hill Publishing Co.), ISBN 0-07-017606-X, 1-800-2MCGRAW.*

THE FIRST INDEPENDENT INTERNATIONAL DISASTER RECOVERY SYMPOSIUM & EXHIBITION SPONSORED BY THE DISASTER RECOVERY JOURNAL & THE DISASTER RECOVERY INSTITUTE

SEPTEMBER 11, 12, 13, 1989

RADISON INN AND CONFERENCE CENTER, ATLANTA, GEORGIA

ABOUT THE CONFERENCE

This will be the largest Disaster Recovery Symposium & Exhibition to date, that is dedicated to disaster recovery. Over 100 vendors are expected. Vendors' exhibits will be state-of-the-art, including major suppliers of services to the Disaster Recovery Industry. We are also expecting four companies that supply mobile services, i.e. a mobile Hot-site/Cold-site. You will be able to walk through and inspect the mobile units. The conference gives you the opportunity to be involved first hand, with hundred of other attendees, in actual case studies. The conference will also offer you credits towards the **CDRP** (Certified Disaster Recovery Planner), **CPE** (Continuing Professional Education) and **CEU** (Continuing Educational Units). You will be exposed to experts in the field who will explain how they complete their tasks of analysis and planning. Most importantly, you will be involved!

GENERAL SESSIONS

General Sessions will feature nationally known authorities speaking on current disaster recovery subjects of interest to all attendees. They come from varied industries and geographic locations. The attendees will have an opportunity to meet the speakers and discuss their own situations. This could lead to problem resolution, without costly mistakes and time delays.

BREAKOUT SESSIONS

There will be an extensive selection of lecture sessions, workshops, and panel discussions to choose from. The sessions will allow attendees to be involved in lectures, hear speakers solving problems, and walking through case studies that relate to their own personal and corporate needs. The panel discussions provide valuable insight into current problems, and focus on tried and proven solutions.

WHO SHOULD ATTEND

The key to corporate survival is disaster recovery planning. Anyone responsible for developing the corporate contingency plan will find the conference an invaluable asset and will assist them in their final goal.

Auditors will get the necessary exposure to disaster recovery planning to help them resolve problems for their clients, and help them in auditing corporate disaster recovery plans.

Directors will find the conference will help them in selling senior management on business survival alternatives.

WE ARE EXPECTING OVER 100 VENDORS. A FEW OF THEM ARE AT&T, HOTSITE, IBM, SUNGARD, WANG, US WEST COMMUNICATIONS, COMDISCO. WE CURRENTLY HAVE 53 VENDORS REGISTERED.

WE ARE EXPECTING 200 TO 300 ATTENDEES.

PARTIAL LIST OF KEY SPEAKERS

G. EDWARD MATTINGLY
NATIONAL BANK EXAMINER EDP ACTIVITIES DIVISION
CONTROLLER OF THE CURRENCY

THOMAS C. JUDD
DIRECTOR OF THE CONTINGENCY PROCESSING CENTER (CULPEPER, VA)
FEDERAL RESERVE BANK

COLE EMERSON
PREVIOUSLY WITH FIRST INTERSTATE BANK, LOS ANGELES
SRI INTERNATIONAL

MELVYN MUSSON
ASSOCIATE VICE PRESIDENT MM & PROTECTION CONSULTANTS
MARSH & MCLENNAN

JAMES BOYLE
VICE PRESIDENT
IBM CORPORATION, NSD

NORM HARRIS
PRESIDENT
HARRIS DEVLIN ASSOCIATES

PAT WILLIAMS MOORE
NATIONAL EDUCATION COORDINATOR
BMS CATASTROPHE

LLOYD R. SMITH
DIRECTOR OF COMMUNICATIONS-COMPUTER SYSTEMS
OKLAHOMA CITY AIR LOGISTICS CENTER

CHARLES E. PERKINS
SENIOR MANAGER
COOPERS & LYBRAND

CALL NOW FOR YOUR BROCHURE! DISASTER RECOVERY JOURNAL (314) 846-1001 – FAX (314) 846-1023

CIRCLE #230 on Reader Service Card ▲

Testing In VM

By Michael Seadle, Ph.D.

How application development teams can use VM to improve their productivity is the subject of this article. It also illustrates how to use VM to test newly developed programs.

This is the second article in a series dealing with VM in the development center. The first, "VM In The Development Center," in the June 1989 issue of *MAINFRAME JOURNAL* presented an overview of development issues and focused on how to make effective use of VM's electronic mail facility.

VM offers two alternatives for testing non-CMS programs: emulation or transfer to a guest operating system. Final testing should always take place in the environment the production version will use. Not doing so might give some unpleasant surprises. Nonetheless, emulating OS or DOS as a first phase of testing lets team members take advantage of VM's debugging facilities.

Programs team members assemble or compile in CMS result in a TEXT deck that is unlinked object code. They must LOAD any other programs necessary for execution, then BUILD the module. The process mirrors what they do when they link a module in VSE or MVS. They just LOAD the main component of the module, INCLUDE any subprograms, then generate the module with a GENMOD. You can build a simple EXEC that can automate the process for them. IBM's ASMGEND EXEC provides an example to work from. If the programs are using OS macros, you should also access the disk the MACLIB is on, point to the MACLIB with a FILEDEF and issue a GLOBAL command to make the macros available.

Using VSAM In VM

VM offers a good environment for testing batch programs, but team members may need help setting up their test files. Programs using a database management system require the VM version. Without it, they cannot test. Ordinary VSAM files cause less trouble. A duplicate VSAM license costs little and installs relatively quickly. You should note, however, that VM uses VSE VSAM meaning you must invoke DOS emulation routines before starting.

You can put a VM VSAM catalog on any minidisk, even TDISK, but it cannot share the minidisk with CMS files. The reason is that the disk must have CP, not CMS formatting. You can use IPL FMT from the S disk to reformat the minidisk, but remember that it will erase any files you have there. You have probably used IPL FMT to configure disk packs or to allocate the bit-map that specifies cylinder usage, but non-systems programmers may not be familiar with it. To help them, you should set up an EXEC that uses IPL FMT's batch control option. If, for example, you need to format a five-cylinder minidisk called TMP01 at virtual address 19F on an IBM 3380, the batch command is:

```
FORMAT,19F,3380,TMP01,000,004
```

You put this command in a file called FORMAT DATA, then stack FORMAT

FIGURE 1

```
/* REXX EXEC to format a disk for a VSAM catalog */
100.MAINLINE:
  /* Direct your PUNCH to an unused RDR class */
  SPOOL PUNCH '*' CLASS P CONT
  /* Put IPL FMT and its batch commands into your RDR */
  'PUNCH IPL FMT * (NOH'
  'PUNCH FORMAT DATA A (NOH'
  SPOOL PUNCH '*' CLASS P CLOSE
  /* Make sure IPL FMT heads the RDR queue */
  'ORDER RDR CLASS P'
  /* Invoke IPL FMT by IPLing your RDR
  'IPL 00C'
  EXIT
```

DATA behind IPL FMT in the reader queue and IPL the reader as in Figure 1. After formatting the minidisk, they can set up a VSAM catalog using standard Access Method Services commands.

To use files in the catalog, team members must establish DOS labels by doing a: SET DOS ON and coding DLBLs. VM VSAM does not use FILEDEFs. They will need a DLBL for the master catalog as well as ones for any user catalogs. VM allows any number of master catalogs, as long as each virtual machine uses only one. Some team members may prefer to let their master catalog serve as a user catalog too. In a test environment, this causes no problems. After the DLBLs are ready, they should SET DOS OFF unless the program requires DOS emulation.

Then begin the test by entering the name of the module. Testing in an interactive environment will not degrade VM the way it does CICS or TSO. This is because the VM scheduler automatically treats a virtual machine that exhausts its interactive Q1 timeslices as a batch Q2 user.

Debugging

If an abend is encountered, examine the dump on-line with the DISPLAY command. DISPLAY shows any portion of virtual memory in hexadecimal format and translates the hexadecimal code exactly as would a paper dump. DISPLAY G will also show the general registers. In order to save the information for later reference, first:

```
SPOOL CONSOLE * START
```

Then everything displayed on the screen will go into a reader file. To stop the recording, enter:

```
SPOOL CONSOLE STOP
```

And to close the spool file and send it to your reader:

```
SPOOL CONSOLE CLOSE
```

Either print the results or RECEIVE them onto the minidisk. The latter is better because it allows the use of REXX and XEDIT to help find the problem.

You may also take advantage of other VM debugging tools such as PER and TRACE. PER will trace Assembler instructions, branches or register changes for particular ranges of program addresses. TRACE displays all virtual machine activity, including supervisor calls and I/O. These commands may be used with COBOL or any other high level language, as long as the compiler's Assembler listing is available as a guide.

At present, VM has no test environment for CICS programs. IBM has announced a CMS version of CICS. However, until this product has more public use, it is difficult to say whether it can provide a sufficiently standard CICS environment to rely on for testing.

Testing In A Guest

Program development in VM does not require testing in VM. Batch jobs may easily be submitted into MVS or VSE guest machines. Both JES and POWER can route the output back to the VM reader queue. Job streams are handled the same way regardless of whether they came from internal or external sources. The biggest difference is that in VM only your own listings are in the queue. You need not search for them.

You can write REXX EXECs to build JCL for batch submission. A compile job, for example, generally needs only a few

FIGURE 2

```
300.SPACE.CHECK:
  MKBUF                    /* Get separate buffer */
  QUERY DISK A '(STACK'
  PARSE PULL LINE /* Discard headings line */
  PARSE PULL LINE /* This is the line you want */
  PARSE VAR LINE . . . . . BLKSIZE . . BLKLEFT .
  /* Blocksize makes a big difference. Check that before
     deciding on space availability. */
  IF BLKSIZE = 4096 & BLKLEFT < 50 THEN DO
    SAY 'TOO LITTLE SPACE AVAILABLE —' BLKLEFT 'BLOCKS'
    EXIT 50      /* Return Code 50 = have < 50 blocks */
  END
  IF BLKSIZE = 1024 & BLKLEFT < 200 THEN DO
    SAY 'TOO LITTLE SPACE AVAILABLE —' BLKLEFT 'BLOCKS'
    EXIT 200     /* Return Code 200 = have < 200 blocks */
  END
  DROPBUF                  /* Drop separate buffer */
RETURN
```

pieces of variable information: the program name, account numbers and other modules for the linkage editor. The EXEC can prompt for these variables, build the job stream in a temporary file (using EXECIO) and COPY the program to it. Then it can submit the job and delete the temporary file.

The most common problem your submit EXEC will encounter is a lack of file space. You can build in a simple routine to check for space given in Figure 2.

Of course you will have to decide how many blocks the temporary file is likely to need. This will depend on your disk type and blocking, as well as the maximum size of the programs. If you have chronically limited disk space, you can build the EXEC to define TDISK, format it (just stack the format commands) and use that space to build your job stream. Remember, however, that the heavy I/O involved in formatting your TDISK can make job submission rather slow.

VM does lack one popular feature present in both MVS and VSE. Team members cannot easily monitor the input queues to discover the relative position and priority of their jobs. Some people check constantly. They set up the commands on a PF key that they hit constantly for long periods to watch their job run. Any programmer who is not a saint has undoubtedly done this. Few will argue that such monitoring does any good. VM removes temptation saving time and perhaps a few ulcers.

On-line testing almost always requires changing from one interactive environment to another. Anyone working in TSO must sign onto CICS and those already in CICS usually need to switch to a different CICS. Likewise in VM, team members must switch in their case to a guest system. VM VTAM should connect them directly to the right region or partition in the guest. Or, if your shop has no VM VTAM, it can DIAL to the guest. You may need to give them a specific address to get the right CICS.

Test Verficiation

Verifying test results is one of the most tedious and important parts of system development. The fact that a series of programs work without abending does not guarantee the validity of their answers. Yet people often fail to check results carefully or neglect to recheck them only to discover that changing an unrelated routine has altered the results. VM can help automate a project team's verification process.

Once team members have a test plan, they need to establish a test database. This can be either a random sample of production records, if available, or manufactured data designed to cover all known conditions and errors. The latter makes a good first-stage test but is no substitute for the unpredictable variety of live data. Team members should keep two copies of this test database: one in files owned by the guest operating system and another in VM in the format used for the system prototype. That prototype can calculate answers to check against the test system reports.

For on-line applications, they can check screen prints from the prototype against the same queries in CICS. Some results may differ because of errors in the pro-

See VM *page 93*

An Experience In MVS Batch Tuning

By Patrick L. Gaul and Radik A. Gens

This article describes an experience in tuning an MVS batch workload at Canadian Airlines International Ltd. The definitions of test job input classes were redefined based on MVS System Resource Manager (SRM) service unit consumption. Installation Performance Specification (IPS) parameters were specified to serve this redefinition. An SMF user exit, IEFACTRT, was employed to provide informational messages for users regarding the performance of their jobs and the users were briefed on the new definitions. With their approval, the implementation of these changes proceeded. The final result was a dramatic improvement in turnaround times for test batch jobs.

In March 1987, Pacific Western Airlines Ltd. and Canadian Pacific Airlines, Ltd. merged to form Canadian Airlines International Ltd. After several months of preparation, the data centers of the two companies were merged into one and the workloads of the two companies combined.

Soon afterward, there was a perception that batch job turnaround time had become unacceptably long. This was noticed in particular for jobs submitted by applications development staff and end users. The combined volume of non-production jobs had affected the performance of this component of the system workload. This group of jobs will be referred to as test batch.

It was time to review job input class definitions for the test batch subsystem and investigate the use of the MVS System Resource Manager (SRM) parameters which affect its performance.

Batch jobs had long been classified according to a criteria established years earlier based on CPU (TCB + SRB) time as shown in Figure 1.

This criteria was established at a time when the MVS workload was run on an IBM 370/158 processor but had never been updated through various processor changes. Naturally, as the MVS processor speed increased through several hardware upgrades, the criteria for Class I jobs became less trivial. By some estimates, 30 seconds of CPU time on a 158 is equivalent to 1.7 seconds on an IBM 3090/180E, which is the current MVS processor.

We knew that this criteria was an outdated one and that we had to begin the tuning exercise by examining the phases in the life of a batch job to see where the bottlenecks lay.

The turnaround profiles of these classes of batch jobs were analyzed for two periods which were judged to be typical; they were each one month long and were one year apart. Our intent was to look at how each phase had changed over the course of a year. We investigated 90th percentile values for the converter, input queue, dataset enqueue, allocation, execution and total turnaround times.

Turnaround time is defined here as the time from job read-in to job termination. Output queue time has been excluded from this exercise, since it would involve many factors which we were not in a position to tune.

For non-statisticians, percentiles are values below which a given percentage of measurements may be found. For example, if the 75th percentile of input queue time is 3,000 seconds, this means that 75 percent of all observed jobs were in the input queue for 3,000 seconds or less.

Figure 2 shows data from a one-month period. Clearly, input queue time was overwhelmingly the largest component of job turnaround time.

When a large number of jobs are waiting in the input queue, this tends to suggest that not enough initiators are available. We decided that, rather than have the jobs waiting in the input queue where we cannot improve their performance, we should get them initiated and deal with their performance through the use of the MVS SRM. So, we added more initiators for each class.

Immediately, the bottleneck shifted. Soon we observed that the CPU was running at a utilization of 101 percent. (Note: the CPU was actually running at 100 percent; the extra percent indicated that the SRM had a task swapped in and ready for dispatching but had not been able to dispatch it.) Now, instead of a shortage of initiators, we were experiencing a shortage of CPU cycles. Moreover, we found that the initiators became occupied by long-running jobs and eventually input queues grew back to their previous levels. We were not getting enough of the jobs through the system and short-duration jobs were waiting long times due to the long-running jobs.

It was time to employ the SRM to solve the problem.

MVS System Resource Manager (SRM)

The SRM is described in great detail in the IBM document *MVS/Extended Architecture System Programming Library Initialization & Tuning (GC28-1149)*. It is

FIGURE 1

Job Input Class Definition

- **Class I (short jobs): less than 30 seconds**
- **Class A (medium jobs): less than 3 minutes**
- **Class F (long jobs): no limit**

not our intention to describe the SRM's full function here. We will just limit our discussion to a few relevant areas.

The main function of the SRM is to distribute the available resources among the various tasks running on the system in order to maximize throughput. The resources which the SRM is primarily concerned with are CPU (TCB and SRB), I/O and central memory.

All SRM resources are quantified in service units. CPU service units are dependent on the processor speed, I/O units on the number of I/Os executed and central memory service units on the amount of central memory used while the CPU is executing. Each type of service unit may be assigned a relative weight through the IPS member IEAIPSxx of SYS1. PARMLIB, depending on which resources require the most balancing.

The SRM can distribute resources through two basic means: swapping tasks in and out of central memory and establishing the dispatching priority of those tasks which are swapped in.

The SRM decides on swapping and dispatching according to performance group definitions as specified in the IPS. Performance group definitions should be made for each component of the MVS workload which may require specific performance characteristics.

Within each performance group, periods may be defined which represent progressive levels of service unit consumption. For example, TSO Period 1 is usually defined for short transactions, using a relatively low number of service units. As the transaction accumulates service, it will reach a limit defined for the period, then fall down into the next period where its swappability and dispatching priority may be changed. It has now become a *medium* transaction. When it reaches the limit of this second period, it will be reassigned to the third period and can be considered a *long* transaction with new swapping and dispatching characteristics.

The MVS workload may be divided into a number of basic workload components. This division must encompass *all* tasks running under MVS.

FIGURE 2

90th Percentile Values Of Batch Execution Phases

	CLASS		
	A	F	I
	90th Percentile (seconds)	90th Percentile (seconds)	90th Percentile (seconds)
PERIOD			
1 Converter	0.97	1.01	0.72
2 Input Queue	76,957.50	102,106.69	428.23
3 Dataset Enqueue	0.15	0.16	0.17
4 Allocation	20.14	49.78	6.31
5 Execution Duration	1,348.48	5,916.51	38.41
6 Turnaround Time	78,232.81	104,421.38	509.21

We categorized our workloads according to function, placing on-line systems (as usual) above batch systems in terms of dispatchability and swappability. The real challenge was to distribute the resources left over after satisfying the critical on-line systems among the batch workloads. Figure 3 illustrates the hierarchy of workloads which were required.

Class O was defined for on-demand, high priority jobs for use by Computer Operations only. Class X was defined as production batch. We will limit our discussion to classes I, A and F.

Since the problem we were experiencing concerned swapping, we concentrated primarily on the swapping attributes of each performance group. Both swapping and dispatching were defined in accordance with each group's relative importance; therefore, in the rest of the article the same principles described for swapping will apply to dispatching priority.

"Know Your Loved Ones . . ."

The problem we were experiencing was clearly that resources were not being distributed among the available jobs properly. We had a gut feeling that we could

FIGURE 3

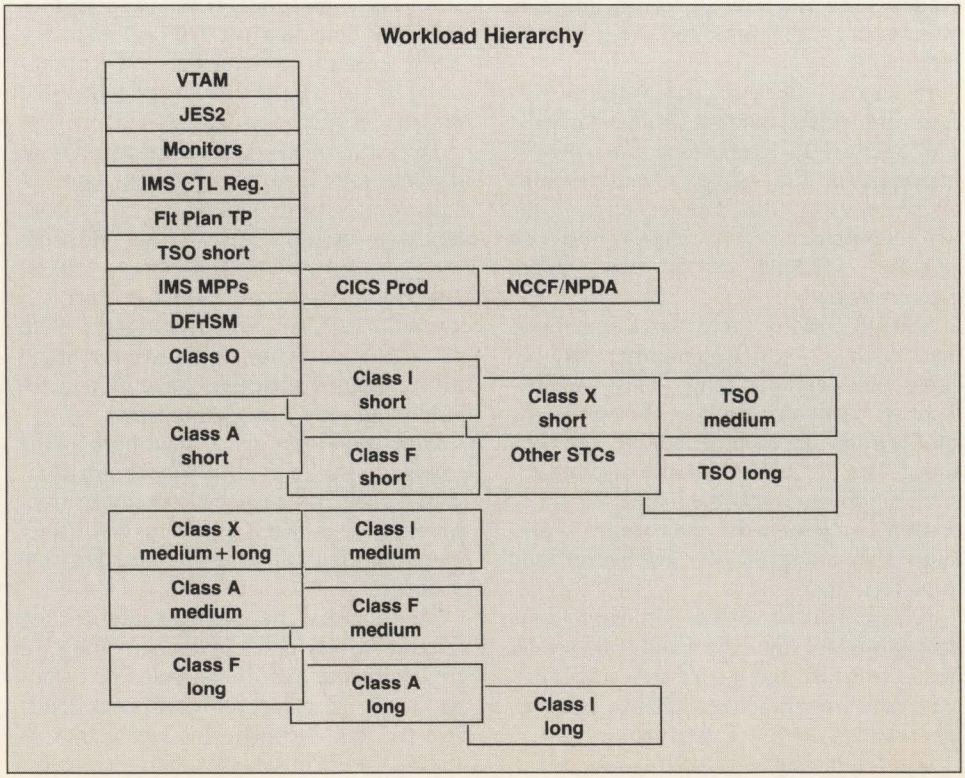

get our short-running jobs through in a fraction of the time that it was taking the long-running jobs to complete. Of course, any changes we made to improve the short jobs would impact the long ones; you cannot get something for nothing. Nonetheless, we felt that a large improvement in turnaround for short-running jobs (which represented the vast majority of total job executions) was worth a modest penalty to the long-running jobs.

We also considered the perception of the user when he submits a job. If he submits a short job, he awaits its completion almost as if it were an on-line transaction; that is, he will suspend further activity until the job has completed. However, should he submit a job that he knows to be long-running, he is more likely to turn to other tasks while the job is executing. In the first instance, he may become quite annoyed if the short job runs for a few extra minutes, especially when he knows that it might run in just a few seconds on an unconstrained system. In the second instance, he might not mind if his job takes an extra hour, if he normally expects it to run for several hours.

The principle of "know who your loved ones are and always have someone else to kick around" ("MVS Performance Legends" by Stephen L. Samson, November 1988, *MAINFRAME JOURNAL)* applies here. The concept is to take service from those components of your workload which can afford performance degradation and give it to those that cannot, whatever the reasons may be (that is, technical, service level agreements and so on).

We decided to rank the priority of batch jobs as inversely proportional to their duration in SRM service units. This was intuitively satisfying, since we discovered that a large number of the jobs were short and by giving them priority, we would increase the total number of jobs through the system in a given time. Also, the jobs would not spend long times swapped out, tying up initiators. We would end up with lower input queues for the short jobs and have more initiators free and available at any given time.

We also decided to ensure that the SRM was balancing the workloads on the basis of *all* types of service units, not just CPU. We knew that although we were experiencing a shortage of CPU cycles, in the future we would inevitably shift the bottleneck to central memory or our I/O subsystems.

The real challenge was to decide how to subdivide the jobs into short, medium and long categories and in a way that the user community could relate to easily.

The Analysis

We began our analysis by looking at the breakdown of service units consumed by the jobs in each class. Figure 4 shows the percentiles.

Examining the data, we came up with the following period limits:

1. Period 1:
   - (All classes short) . . . . . 10K units
2. Period 2:
   - Class I medium . . . . . 50K units
   - Class A medium . . . . . 1,500K units
   - Class F medium . . . . . 500K units

The limits for Period 2 for each of Classes I and A were chosen in accordance with 90th percentile values. The limit for Class F Period 2 was made lower than the 50th percentile since Class F jobs were typically longer-running, yet we wished to retain a modest intermediate range while giving preference in this range to Class A jobs. Period 3 for Class F would have

FIGURE 4

Service Unit Percentiles By Job Input Class

	CLASS					
	A		F		I	
	Service Units	Jobs	Service Units	Jobs	Service Units	Jobs
Percentile						
1st	470	19	1,517	7	82	469
5th	583	93	3,105	33	293	2,345
10th	2,753	186	11,742	67	333	4,690
25th	18,285	465	57,993	166	457	11,725
50th (Median)	145,274	930	931,950	333	855	23,451
75th	666,512	1,395	2,708,542	499	7,893	35,176
90th	1,138,744	1,674	8,733,043	599	38,588	42,211
95th	1,599,578	1,767	18,398,096	632	65,467	44,556
99th	2,839,155	1,841	33,943,904	658	251,402	46,432

the highest priority of the long-running jobs, compared with Period 3 for Classes A and I.

Depending on a job's final service unit total, if it were a Class I job and it exceeded the Period 2 limit, it would be recommended for Class A or F. If it were a Class A job and went to Period 3, it would be recommended for Class F. However, in looking at the number of jobs above the Period 2 limit for Class I (14,000 in a month), we were concerned that this would meet with considerable dissatisfaction among the users, to say the least. So, we backed off and defined a limit of 200,000 units for this class.

Clearly, this was a redefinition of job input classes, based on SRM service units rather than CPU time. The next major obstacle was that the average user did not know an SRM service unit from an isoenzyme.

It was time to talk to them.

The Crusade

We had to find a way to familiarize the users with the concept of SRM service units in a way they could understand. We also had to help them make decisions on which input classes would be most appropriate for their jobs; the definitions would ensure that *if the job was submitted to the most appropriate class, it would receive the best turnaround possible.*

We began by coding an SMF exit, IEFACTRT, that would write messages to the job log for each job indicating the number of service units consumed in each job step. It would also print a message at the end of each job that would recommend the job should have been run in another class, if that had been the case. Figure 5 shows an example of a sample output from the exit.

Then, just before implementing the use of this exit, we met with representatives of the various user groups concerned. In these meetings we outlined the problems and presented the concept of the SRM service unit. We showed examples of the messages that the new exit could issue and how users could apply the information to improve the turnaround of their jobs.

We also told the users that for the first few weeks, we would implement an automatic procedure of compiling lists of jobs that were run in inappropriate input classes and we would send this list out to each user via electronic mail on a daily basis.

It was explained to the users that if they used the definitions correctly, their jobs would run with the best performance available; if Class I jobs exceeded their Period 2 thresholds, their performance could be severely degraded. We also explained that in the latter instance, *the console operator would have the option of cancelling jobs which were holding up initiators for an excessive period.*

We received agreement from each of the user groups and went ahead with our new SRM definitions and the new exit.

SRM Definitions

Performance group definitions were created for each job input class (I, A and F). Within each group, three periods were defined for short, medium and long jobs.

FIGURE 5

Sample Output From IEFACTRT Exit

```
IEF097I GAULPLX  — USER GAULPL      ASSIGNED
ICH70001I GAULPL    LAST ACCESS AT 16:18:29 ON THURSDAY, APRIL 27, 1988
$HASP373 GAULPLX  STARTED — INIT  9 — CLASS A — SYS PMVS
IEF403I GAULPLX  — STARTED
>Step: MXG        .ANALYZE RC=00 Serv=4431K EXCP=4497
>                          CPU=2.20 SRB=.00 Swaps=0 Elaps.time=8.4
IEF404I GAULPLX — ENDED
>=============================================================
>              Tape     Pages     Service      I/O      Total      Total     SWAP    Elapsed
>            Mounts    In/Out      Units    Count    CPU min    SRB min    Count  Time min
>Job Totals:      0       377      4431K     4497       2.20        .00        0       8.4
>
> NOTE:     This job should run in class 'F'.
>
>=============================================================
$HASP395 GAULPLX  ENDED
```

FIGURE 6

Performance Hierarchy For Test Batch

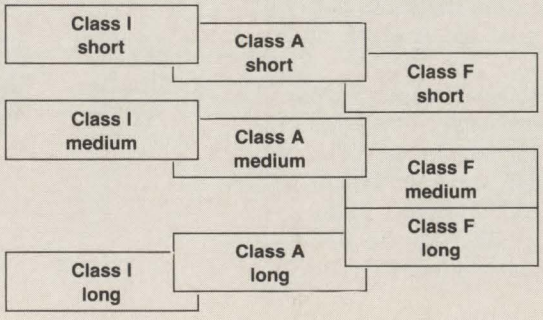

New for your MVS or VM environment!

Are you under the gun to create documentation?

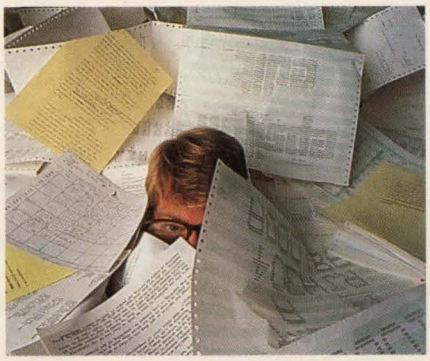

HERE'S AN EASY WAY TO BITE THE BULLET.

Everyone knows that their Data Center should document important policies, standards, and procedures. The trouble is, with all the crushing pressure of work, documentation always seems to go on the "back burner." These days, that can land you in hot water. With all the strict regulations and tough compliance requirements you face, it becomes *absolutely necessary to take action* and finally solve the "documentation dilemma."

Luckily, we can make compliance easier than you ever dreamed possible!

Introducing the Information Systems Series.™ High caliber documentation the painless way.

The ***Information Systems Series*** documentation framework is designed to help you create comprehensive policies, standards, procedures, and user support materials for your entire MVS or VM Data Center.

Created by highly experienced data processing experts, the ***Information Systems Series*** is an interlinked set of five manuals and accompanying, corresponding diskettes: *The IS Policy Guide, The Data Center User Guide, Systems and Programming Standards, The Operations Guide,* and *The Project Plan and Implementation Guide.* All you have to do is follow the simple, imbedded instructions. In other words, you *personalize* the documentation using your own IBM PC or compatible – which means you can easily integrate any existing documentation quickly and easily!

An important part of the ***Information Systems Series*** is our unique *Project Plan and Implementation Guide* that leads you effortlessly through the entire documentation process. No wonder many forward-looking government, business, banking, medical, and telecommunications IS Departments have put the power of the ***Information Systems Series*** to work.

FREE Demo Disk!

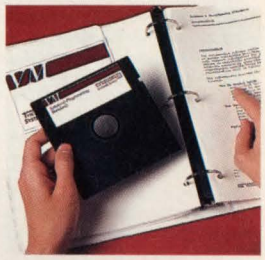

See for yourself what the ***Information Systems Series*** can do for *you.* Call for your free Demo Disk right now! Isn't it worth a few moments of your time to prove to yourself that the ***Information Systems Series*** can save you endless hours of work and many thousands of dollars? Don't miss this no-obligation offer. Call right now!

Call toll-free 1-800-777-0970

The Computer Resources Group, Inc.
303 Sacramento Street
San Francisco, CA 94111

CIRCLE #69 on Reader Service Card ▲

The hierarchy of these periods is illustrated in Figure 6.

The durations of the short periods for all classes are all the same: 10,000 service units. Through the use of performance objective specifications, the swapping characteristics for each class were defined to create the hierarchy shown in Figure 6. Quantitatively, here is a description of the structure defined.

- If a job consumes less than 200K units, Class I is the recommended class. Non-class I jobs will move to Period 2 at 10K units and will experience performance that is inferior to Class I Period 2.
- If a job consumes more than 200K but less than 1,500K units, Class A is the recommended class. Class I jobs will move to Period 3 at 200K units and will experience performance that is inferior to Class A Period 2. Class A will still be better than Class F.
- If a job consumes more than 1,500K units, Class F is the recommended class. Class A jobs will move to Period 3 at 1,500K units and will experience performance that is inferior to Class F Period 3.

FIGURE 7

New 90th Percentile Values For Job Turnaround

	CLASS		
	A	F	I
	90th Percentile (seconds)	90th Percentile (seconds)	90th Percentile (seconds)
PERIOD			
1 Converter	1.12	1.05	0.86
2 Input Queue	41,507.09	62,914.59	209.45
3 Dataset Enqueue	0.28	0.24	0.27
4 Allocation	25.03	36.74	7.37
5 Execution Duration	2,286.49	9,474.48	43.82
6 Turnaround Time	42,177.84	69,859.25	372.07

The IEFACTRT exit would issue messages to jobs which were submitted in a class that would receive poorer service from the SRM than the final service unit count warranted, for example:

- A job submitted in Class I which ended with more than 200K units (that is, Period 3) would be issued a message that the job should be run in Class A or F, depending on the final total.
- A job submitted in Classes A or F which ended with less than 200K units would be issued a message recommending Class I for better performance.

The Results

Figure 7 shows the results of the changes. Input queue time was reduced dramatically for all three classes. This re-

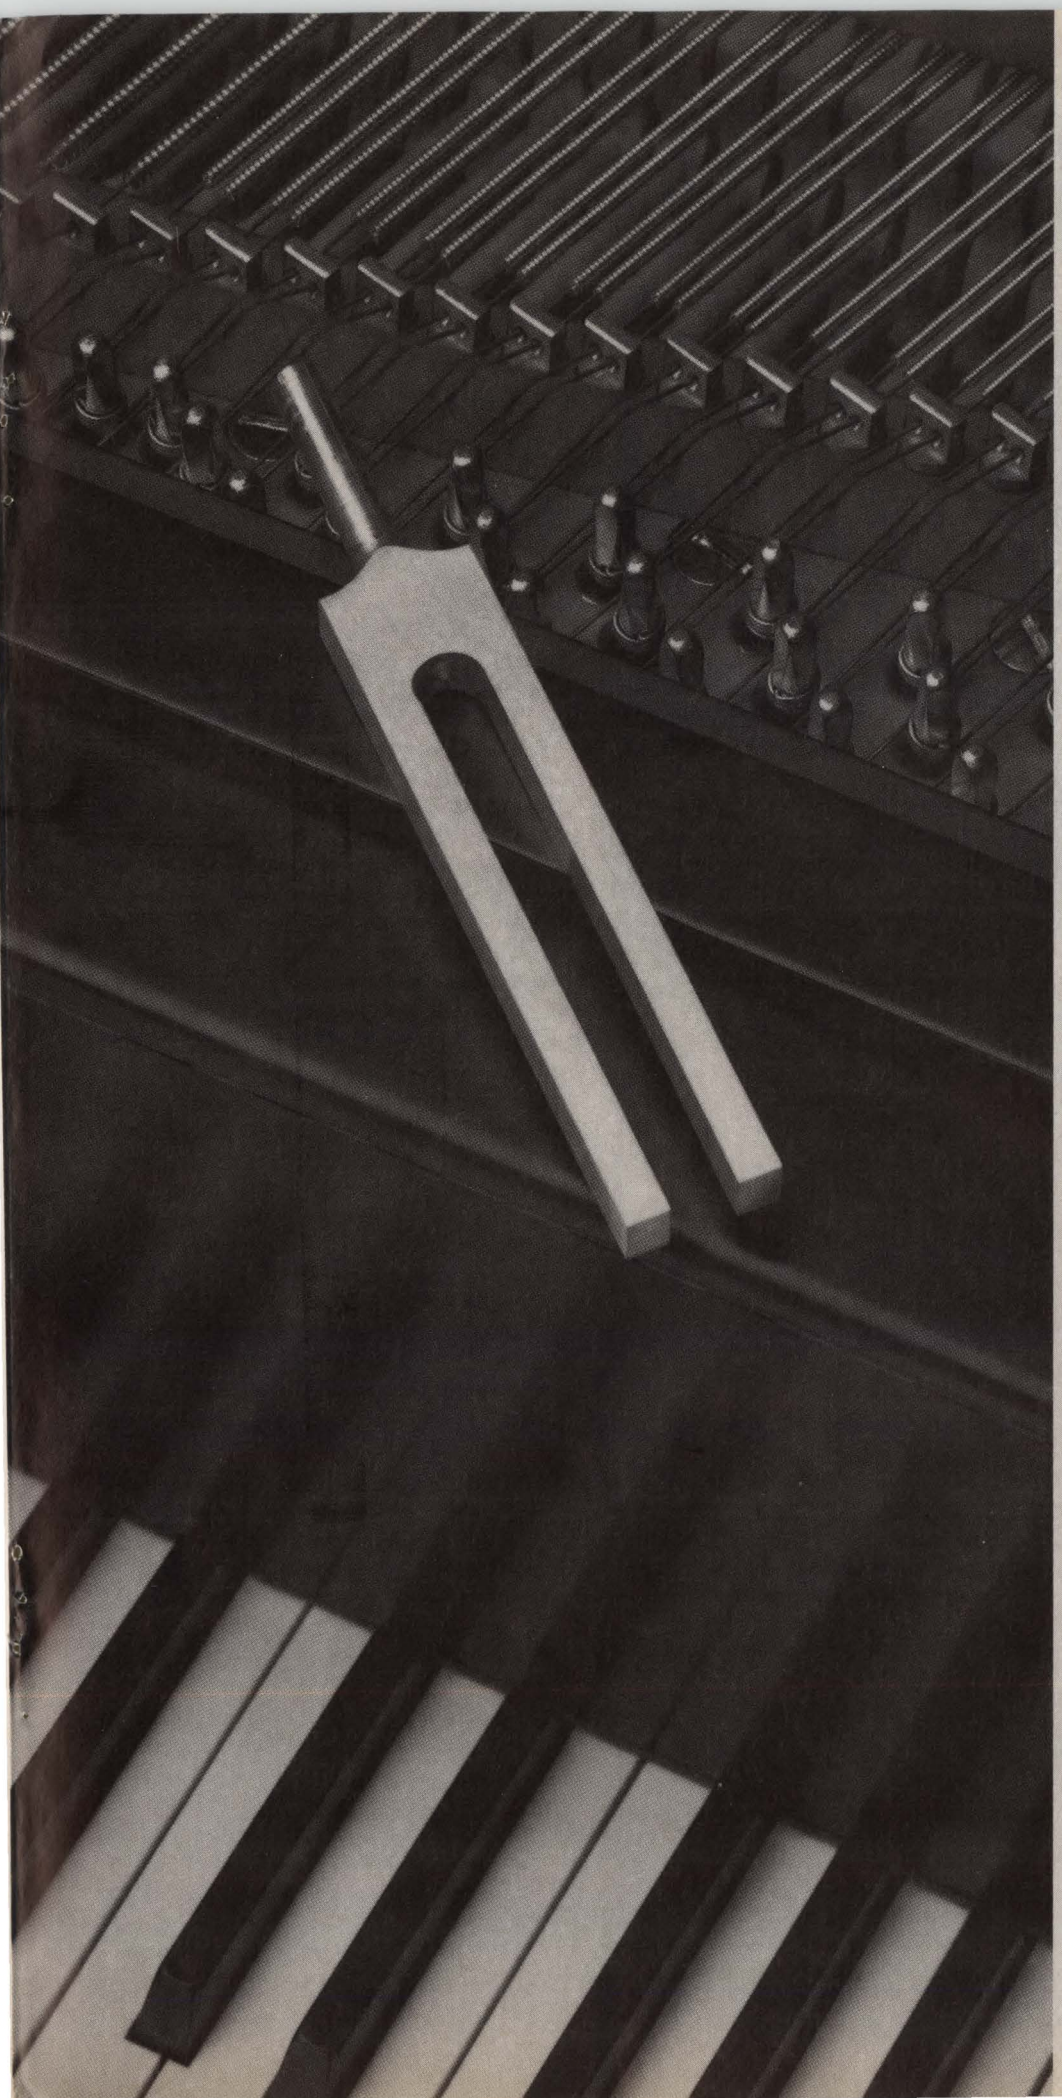

STROBE . . . The Key To Well-Tuned Applications

You know your peformance is great when:

- your customers expect and get fast online response time
- your management reports regularly meet deadlines
- you confidently schedule batch processing before online start-up
- you can bring up new applications without straining your resources

Even if you have already tuned your MVS system, you can assure greater performance with STROBE, the premier application tuning product. STROBE will:

- furnish detail not available from system tuning software
- attribute resource use to specific causes within application programs and subsystems in any address space
- show CPU use by source program statement or system service function, and show disk unit access time by cylinder within data set or index

STROBE measures and reports on programs written in COBOL, PL/1, Fortran, and Assembler, operating in all MVS system environments, and in CICS, IMS, DB2, IDMS, and other subsystem environments.

CIRCLE #33 on Reader Service Card ▲

PROGRAMART
1280 MASSACHUSETTS AVENUE
CAMBRIDGE, MA 02138
617-661-3020 FAX: 617-864-6558

P · R · O · G · R · A · M · A · R · T

sulted not only from the increased number of initiators available, but also because of an increase in throughput.

Input queues could be further reduced by adding more initiators, but in our JES2 environment this could lead to some operational difficulties. For example, a limited number of tape drives may be available and jobs may simply wait in allocation for unit availability; or data-set enqueue wait time could become significant.

Turnaround time was also reduced. This also included a significant level of latent demand jobs; the number of jobs submitted in a month actually increased by about 20 percent, yet the reduction in turnaround was sustained.

It is interesting to note that execution time increased. As more jobs were initiated, the SRM had more tasks among which to share resources, so each job tended to remain resident for longer, awaiting SRM service.

Conclusion

This exercise gave us a great opportunity to examine the power and flexibility of the SRM in managing the workload. The following are key lessons that we learned.

- Make as many initiators available as possible. The IPS gives the performance specialist tremendous power for improving batch throughput, but there is not a thing that the SRM can do about jobs that are waiting in the input queue. Overinitiation is not as bad as it sounds.
- "Know your loved ones." Identify the jobs that you can get the most benefit from by improving their performance whether the criteria be response time, turnaround time or throughput. Get the biggest bang for the buck!
- Get the users involved. Some of these issues simply *must* have the understanding and approval of the user community; for example, job input class definitions. You cannot do it in isolation, so do not try. Give the users a clear understanding of what you are doing and stress the fact that you are trying to help them.
- Watch for the next bottleneck. Invariably, one tuning exercise will end up shifting focus to the next limitation to utilizing the system's full capacity. We saw this happen in this exercise. Do not be discouraged; it is the nature of the game and it is also what you are being paid for.

This was a satisfying experience for us.

See MVS *page 85*

VSE/SP Tips *And* Techniques

IPL, ASI And JCL Commands

By Mark Hanna

The new Interactive Interface with Release 2.1 is similar to the old Interactive Productivity Facility (IPF), but it has many additional functions.

Sometimes you should use the interface all the time and other times, not at all. Switching between using manual methods and the interface can lead to trouble. For example, if you are maintaining the standard label procedures manually and then use the Interactive Interface to create a file or library, the interface will generate a label that will be physically added to the standard label procedure STDLABUP in IJSYSRS.SYSLIB. The next time the manually maintained label procedure in ICCF is submitted, the label added by the interface will be overlayed. So make sure you manually add the same label that was added by the interface to the member in ICCF.

IPL Commands

During IPL, VSE will display a message (1I93I) that gives the percent of space used on the recorder file. When the recorder file becomes full, logging will be stopped. Based on Murphy's Law, the next error that would have been logged will be the information needed to solve a problem. Operations should watch for this message and run EREP to pull off history data and clear the file when it is getting full.

Supervisor Parameters Command

The supervisor parameters command is the first IPL command in the IPL procedure. It does not have an operation field, just operands. The VIO and VPOOL operands of this command will be discussed first. VSE reserves a virtual I/O work area at the upper end of the system GETVIS area and an area of the same size on the page dataset volume immediately behind the last page dataset extent. The area on the page dataset is referred to as the VIO area. VPOOL specifies in kilobytes (K) or megabytes (MB) the size of the page dataset space used to access parts of the VIO space by the LINKAGE EDITOR, POWER and CICS. The VPOOL default size is 64K. A specification of 128K is a good starting place for MODE=E.

With the introduction of VIO, no library or LIBDEF is required for linkage editor output (// OPTION LINK). Virtual storage address space is used instead. Link edit and fetch performance is much better since there is no DASD I/O. Link edit output still goes to a library when "// OPTION CATAL" is used. When a link-and-go is performed, the phase is loaded into the partition from VPOOL. If several link-and-go executions are running at the same time, add additional storage. To run concurrent link edits, VPOOL must be as large as the sum of the number of link edits times 32K for operation in E or VM mode or times 64K for operation in 370 mode. An alternative to a large VPOOL is to use "// OPTION CATAL" and test libraries.

VPOOL storage is taken from the top of all virtual address space. The available virtual storage in each address space is reduced by the size of the VPOOL. VPOOL is regarded as part of the VIO workspace. When using VIO space, the VPOOL area is used for physical access.

The VIO operand (not used for MODE-VM) defines the size of the VIO workspace in kilobytes or megabytes. It must be equal to or larger than the VPOOL operand value. The default is VIO equal VPOOL size. The VIO operand will be the sum of all concurrent VIO users. It should be defined large enough to hold the largest application program but cannot exceed 40MB. Power V2.3 also uses the VIO area to keep a copy of the queue file. The formulas in the POWER installation and operations guide should be used to develop the amount to add to the VIO and VPOOL specifications.

Also on the supervisor command is the VSIZE command used to allocate virtual address space VAE areas. It specifies the maximum total size of all virtual areas which can be allocated in the system. It determines the size of the page dataset.

Parallel page I/O for the page dataset is new for Release 2.1.x and later releases. It allows for overlapping I/O operations supporting multi-extent page datasets on multiple volumes.

Seek time for DASD access should be minimized by dataset placement of the lock file (DLF command), for shared DASD and the label area (DLA command).

The Fast-CCW-Translation (FASTTR) function causes buffers to be saved after an I/O request has been serviced fixing these I/O areas in processor storage until end of job. Using FASTTR will probably increase program working sets since fixed I/O areas will not be released after an

I/O. The VSE supervisor scans its copy/work block buffers for a match. If found, the I/O will be reissued immediately. It works well for programs issuing repetitive I/Os or using large blocks of records. It should be turned off if a job is unlikely to reuse buffers and fixed I/O areas.

CICS and SQL should be executed with FASTTR turned off. FASTTR can be turned off for all jobs by the STDOPT statement "STDOPT FASTTR=NO" in the BG ASI PROC or for specific jobs (such as CICS and SQL) by placing "// OPTION NOFASTTR" in their execution JCL. Since CICS and SQL do not do repetitive I/Os on the same files, using FASTTR can cause performance degradation by causing a shortage of page frames. FASTTR causes I/O areas of user programs to remain fixed in the page pool rather than releasing them after every I/O.

The optional SYS command has three operands that are candidates for tuning. These are BUFSIZE, CHANQ and DASDFP. The BUFSIZE operand is used for MODE=370 and MODE=E supervisors. It allocates the number of supervisor buffers to be used for I/O processing. These buffers (referred to as copy/work blocks) are used by VSE for developing the real addresses of I/O areas from programs that run in virtual storage. If VSE is running MODE=370, then BUFSIZE should equal four times the number of CHANQ entries. For MODE=E, use three times the number of CHANQ entries. BUFSIZE is invalid for MODE=VM. A supervisor generation option that affects BUFSIZE is FASTTR. It is specified in the IBM supplied supervisors. When using FASTTR, the BUFSIZE figure must be increased by approximately 200 percent for MODE=370 and 300 percent for MODE=E. This difference in increase is due to the difference in the buffer size for the two modes. IBM recommends that you specify 1000 initially.

During IPL, Operations should watch for the message that gives the percent of space used on the recorder file.

The CHANQ operand of the SYS IPL command specifies the number of channel queue entries to be allocated and should be set at 255. This should be the default value in the supplied supervisors. Running out of channel queue entries causes performance degradation.

The SYS command operand DASDFP is only needed for DAM files and for processing at the channel program level (PIOCS) for DASD files. It should be specified as NO.

The SVA command is used to allocate storage within the SVA to hold user phases. Having phases in the SVA speeds up processing because it avoids loading phases from disk and it reduces processor storage demands as phases are shared. It is mandatory and must be the last command entered during the IPL procedure. Its operands must be set to match the number and size of phases to be loaded with the SET SDL command in the BG procedure. Place the SET SDL command in a PROC with the phases to be loaded in the BG procedure executed at IPL time. Remember to place a LIBDEF SEARCH in front of the SET SDL procedure so VSE can locate the phases to load into the SVA.

Using the SVA reduces load/fetch time for phases and allows them to be shared across partitions reducing the actual storage size needed for each partition. For example, placing SORT in the SVA allows all partitions to use the same sort phases. They do not have to be loaded into the partition when SORT is executed or called.

Use the following librarian commands to list the IBM supplied SVA load list: ACCESS S=IJSYSRS.SYSLIB followed by a LIST $SVA*.PHASE command. This will be a listing of the names of the phases that the VSE system loads into the SVA at startup.

When cataloging SVA-eligible phases, always check the LNKEDT listings to see if they were loaded successfully into the SVA. New versions of SVA-resident phases are dynamically added after any existing phases in the SVA if there is room. If the SVA is full, the phase is cataloged into the library specified but is not loaded into the SVA. Use the LIBR list directory command for the SDL to get your SDL and SVA usage. The output of this command is a listing of the SDL total number of entries, used entries, free entries and their percentage of the total for each; for the SVA, the memory total in kilobytes of total space, used space, free space and the percentage of total space for each.

There is a procedure in the system procedure library called ''SETSDL'' that contains frequently used B- and C-transients phase names and ICCF phase names. The SET SDL command is issued from the SETSDL procedure and in the BG ASI procedure to indicate the phases to be loaded to the system directory list and SVA. In addition to the IBM supplied names running CICS eligible phases in the SVA reduces real and virtual storage requirements when more than one CICS is being executed on a VSE system. The *CICS Installation and Operations Guide* lists the eligible phases, their SVA storage requirements and how they are installed.

ASI Processing

VSE Release 3.1 and later provides five special system startup modes: COLD, BASIC, MINI, WARM and RECOV. WARM and RECOV are reserved for use by VSE/SP. During the IPL process, the operator has 10 seconds to interrupt partition startup processing and request one of the other three types of startup. VSE will automatically perform a WARM startup for each partition if a normal shutdown of all partitions was completed successfully. A RECOV startup is automatically started if shutdown was not successful or no shutdown was performed. Most of the time the startup will not be interrupted. Computer operations should be trained on this feature.

On a COLD start the POWER queues are formatted and selected jobs are loaded into the queue. The jobs to be loaded into the new queue are defined in the COLDJOBS procedure located in IJSYSRS.SYSLIB. A skeleton (SKCOLD) to load the COLDJOBS procedure is in ICCF library 59. A skeleton (SKLOAD) in ICCF library 59 can be used to load the jobs you would want loaded into the new POWER reader queue during a COLD start. The job streams are cataloged into IJSYSRS.SYSLIB with a member type of Z. VTAM and the CICS startups are good candidates. The VSE/SP administration manual covers these startup skeletons. Note that ''* $$'' JECL statements are coded ''$$$$'', ''/*'' is coded ''$$/*'', and ''/&'' is coded ''$$/&''.

The MINI startup procedures should be documented for emergency library backup/restore. Only BG and F1 are activated. In a DASD sharing environment, the MINI startup should be performed regularly so the librarian RELEASE command can be issued for IJSYSRS.SYSLIB to release

deleted library space. Otherwise, a full library condition may occur from VSE writing control information in the library at startup. Also, performance problems may be caused by a large amount of "delayed space" in a library.

The BASIC startup will be used when normal system startup fails or VSE cannot be started normally. The BASIC startup will allow the system to run so system modification errors can be corrected. Both the BASIC and MINI startups should be tested to make sure they are working. Do not wait until they are needed. They might not work after the system has been tailored. You should not perform a BASIC or MINI startup on two VM installed VSE/SP guest machines at the same time.

The PRTY command is used to set default priorities. It should be issued in the BG PROC to specify the priority and balancing for each partition. It may be issued by the operator to dynamically alter partition balancing. Besides assigning a fixed priority to partitions, two or more partitions can be defined to be treated equally in priority. POWER (F1) must always have the highest priority followed by VTAM (F3) and CICS (F2). If all other partitions are to be equal in priority, the command to set this priority would be: PRTY BG=FB=FA=F9=F8=F7=F6=F5=F4,F2,F3,F1. The MSECS command can be used to alter or display the partition balancing time slice.

The TPBAL command may be used to further alter partition balancing. Processing will be delayed in the specified number of partitions of the lowest priority. This command can help when there is excessive paging caused by the CICS partition. It trades batch response time for CICS response time. It can have an unacceptable negative effect on batch processing. TPBAL works best for low activity CICS systems with small real storage allocations.

The ALLOC command allocates virtual storage to partitions. The SIZE command is used to specify the amount of virtual storage to be reserved for a partition. The SIZE command value is subtracted from the ALLOC command value to produce the GETVIS area size. This value should result in at least 128K of GETVIS in each partition.

JCL Considerations

JCL standards should be developed to keep JCL simple and understandable. Comments should be used in the job stream to explain what is happening.

Trying to locate occurrences of file names or other information in JCL streams and programs that are maintained in ICCF can be time consuming. I have updated an IBM Installed User Program (IUP) product for my ICCF clients that IBM does not support for SP/2 and later. It is an ICCF library "SKAN" utility that runs batch or interactively. It has been used to locate and change file names and SYS numbers in JCL and programs. The maintenance task is made easier if standard ICCF libraries are established. Consider using a separate ICCF library for each of the following: production JCL, CICS, VSAM JCL, VTAM and NCP, system maintenance (ASI and IPL procedures), non-IBM product installation JCL, FCBs, documentation, sample utility JCL (such as LIBR and ICCF) and ICCF procedures and macros that are also in the common library.

Cataloged procedures may now invoke other procedures. Catalog procedure nest-

ing is used frequently in the IBM ASI procedures. These nested procedures may be nested 15 levels. A procedure may not call itself or call a procedure from which it is called. All the procedures in a nesting must specify DATA = YES or DATA = NO.

Procedures may also pass symbolic parameters. Symbolic parameters are assigned or modified at execution time. The SETPARM JCL statement defines and assigns a value to a symbolic parameter that may be used in later job control. It is effective for only one job. If a parameter is defined in an EXEC PROC or PROC statement, it is valid until End-of-Procedure. If the procedure is defined by a SETPARM statement from SYSRDR, the parameter is valid until End-of-Job. If the SETPARM is in a procedure its parameter is valid until End-of-Procedure.

Use LOGSRC on the // OPTION statement along with LOG for debugging symbolic parameters. It causes JCL which contains symbolic parameters to be printed twice: once in source form (as it was coded), once with substituted symbolic parameters (as modified by job control). When these two options are in effect, all statements skipped by job control are written to SYSLST along with a message giving the reason it was skipped.

Conditional job control statements allow job-step to job-step communication. Return codes are set by IBM components and may be set by user application programs. This new JCL provides logic to alter job step sequence.

An example of conditional JCL to bypass the LNKEDT step if the maximum return code is greater than zero would be:

```
// IF $MRC GT 0 THEN
// GO TO NOLINK
// EXEC LNKEDT
/. NOLINK
```

There are three label procedures provided for VSE. The STDLABEL procedure contains system standard labels. The STDLABUS procedure is designed to contain user labels for non-VSAM files and partition standard labels. The STDLABUP procedure is designed to hold VSAM file labels including those labels for files created by the Interactive Interface.

I prefer to place all IBM product system files that are not VSAM in STDLABEL, all IBM product files that are VSAM in STDLABUP and all user labels for all file types in STDLABUS. I maintain the label procedures in ICCF instead of using the Interactive Interface.

See VSE/SP *page 85*

Managing The Complexity Of The DB2 Environment

DB2 MASTERMIND offers solutions that help.

By Mary Lou Roberts

In 1985, IBM at last joined the ranks of relational database management system vendors with the introduction of DB2. Not surprisingly, although IBM was a late entry into a market previously dominated by companies like ADR (now a division of Computer Associates) and Cullinet (also a division of Computer Associates), the company quickly grabbed the lion's share of attention and users. This is, after all, IBM.

Again, not surprisingly, the quick growth in the ranks of DB2 users created a whole new after-market for independent software vendors who develop and deliver products designed to reduce the headaches of IBM users.

Since the inception of the proprietary software industry, the IBM after-market has been a mainstay of product revenues. Lots of money has changed hands between users of IBM equipment and independent vendors based on the following statements: IBM does not offer it and we do (for example, CASE products); IBM offers it, but we have a better one (for example, sort utilities); and we offer a product that will help you use/manage/increase the performance of IBM hardware and/or software.

The products of BMC Software fall into this third category. BMC, based in Sugar Land, TX, was founded in 1980. Since that time, it has made its mark quickly and profitably by marketing enhancements to IBM's database and data communication program products. For the first several years of its corporate life, that meant selling to the users of IMS and CICS. Now DB2 has opened up an entirely new window of opportunity and BMC appears to be climbing through that window at top speed.

DB2 MASTERMIND: The Database Administration Series

IBM's DB2 purports to ease the process of application development. That is, of course, one of the major drawing cards of having a relational database management system. But, in doing so, the behind-the-scenes aspects of database administration and maintenance have become far more difficult.

BMC's delightful product brochure, written and illustrated in comic-book fashion, depicts the tentative DB2 user — a quaking database administrator. Entering the land of DB2, he finds himself "standing on the edge of a frightening precipice, looking into a future of challenges and dangers that no DBA should have to face alone." He cries, "I enthusiastically took my first step. But when I looked back, I was alone and in trouble! . . . Our choices have been taken from us. Any changes now to our plan will cause this creature to destroy us!" Enter the character DB2 MASTERMIND who "WHAPs" the dragon and promises to control the "beast" with a single mind wave.

DB2 MASTERMIND leads the user through the paths of DB2 where together they encounter the horrible skeletal remains of DBAs and "an incredible stampede of frantic, wild, angry unrecognizable objects" that "multiply before our eyes." The superhero, with a wave of his hands, controls the savage objects and leads the user to the land of relational database management unscathed.

Of course, the DBA user and DB2 MASTERMIND oversimplify the facts. They are no more representative of real life than are Archie and Veronica. Yet, BMC makes its point well: the DB2 environment is fraught with peril and tools are necessary to manage it and navigate through its plethora of new data structures if one is to survive with any semblance of productivity and order.

To this end, BMC offers DB2 users DB2 MASTERMIND, three administrative software tools, all of which utilize a standard user interface.

DB2 ALTER

While DB2 provides powerful facilities for constructing complex data structures, it offers few tools for changing them. DB2 users may find themselves with increased flexibility for querying their database. But IBM does not extend that flexibility to the process of altering the data structures. DB2 database administrators, then, find themselves with a highly-complex task at hand when they want to make changes to

a table. Often, they are required to rebuild the structure completely — from the ground up.

DB2 ALTER is designed to offer support for changing, copying and migrating DB2 data structures, reducing significantly the work of the DBA. Its primary features include the abilities to:

- Automatically restore data, dependencies and authorizations through global changes made with a single command
- Support a large number of data conversions for both format and data type changes
- Allow users to specify changes to the data structures interactively; these changes can then be executed in a batch mode
- Support changes to all attributes of data structures
- Provide a single command (MIGRATE) which can, in combination with the Database Node and Global Change features change, move and install an entire application's data structure
- Automatically execute required DB2 utilities
- Offer synchronization point control and full restart capability.

In addition to relieving the database administrator of many tedious tasks, DB2 ALTER has two other major benefits. It eliminates change backlogs and ensures the integrity of production applications that have had changes made to them.

DB2 DASD MANAGER

The second product in the series, DB2 DASD MANAGER, is designed to ease the tasks of data management by setting up an object, collecting and analyzing statistics as the object grows and executing utilities that assure an acceptable level of performance.

The operational characteristics of DB2 underline the importance of these tasks. The way an object is set up in DB2 will directly affect system performance. If DB2 default allocations are used, system performance can be seriously degraded. To manage DB2 datasets for minimum DASD volume contention and maximum performance, DBAs must allocate the datasets themselves.

In addition, the statistics function of DB2 leaves something to be desired. It is slow and incomplete. The DBA should have access to better information in order to choose whether or not to change catalog values.

BMC's DB2 DASD MANAGER provides a capacity planning and performance tuning tool for controlling the life cycle of DB2 physical objects. Its primary features include the abilities to:

- Provide comprehensive space management statistics and estimate space and selects volumes
- Generate AMS commands and utility job streams
- Provide a historical database for statistics and utility event recording
- Check utility job streams for object validity
- Eliminate unused allocated DASD space
- Allow the DBA to choose whether or not to change the catalog values.

With DB2 DASD MANAGER, companies can standardize the way in which certain DB2 administration functions are performed across the entire corporation.

DB2 CATALOG MANAGER

DB2 CATALOG MANAGER is designed to increase system productivity by providing quick and easy access to DB2 catalog information. It also lets users with little previous knowledge of DB2 or SQL perform catalog maintenance. Major features include the abilities to:

- Allow execution of utilities, DB2 commands and SQL statements against displayed objects
- Support copying authorizations from one object to another
- Provide reliable recovery of structures
- Deliver an audit trail and print all lists and object descriptions.

Overall, DB2 CATALOG MANAGER is designed to offer increased versatility in the development of DB2 applications and improve system security and integrity.

User Response

Darrel Stewart, Manager of Data Administration at Bear Creek Corp. in Medford, OR testifies to the need for DB2 management tools.

"With DB2 it's easy for the database administrator to become a bottleneck for the applications people and that shouldn't be the case. The amount of information a DB2 analyst has to deal with in this environment is huge. To manage it without help would be a tremendous nightmare," Stewart points out.

Bear Creek has installed all three of BMC's DB2 administrative series of products. "We plan to use the migration tool to migrate DB2 from testing level to testing level. Migration tools give you the ability to take relationships from different testing levels and re-establish correctly the relationships," Stewart explains.

Stewart believes that database administrators working in the DB2 environment simply must have tools that let the computer do a lot of the work for them. "We currently have about 45 developers working on one major DB2 project. The only way that I, as a DBA, can support this number of developers is with help and BMC's products give me that help."

Mark Townsend, the supervising systems programmer in charge of database and data communications products at the City of Los Angeles Department of Water and Power, agrees with this assessment. "We were looking for a product that would ease the burden on DBAs and make simpler the complexity of change that is inherent in the DB2 environment. Making DB2 changes has become so monumental — so complex," he explains.

"We also wanted a product that would help manage our DASD resources. Other products we looked at wouldn't give us the help we wanted below the block level. DB2 DASD MANAGER did," he states.

Both Stewart and Townsend praise the quality of BMC's software, documentation and support. "BMC is constantly looking for ways to improve its products," Stewart says.

Like its forerunner, IMS, shops that commit to DB2 will make a major investment — perhaps larger than they had anticipated. And, just as IMS and CICS spawned their own sub-industries of computer software support products, so will DB2.

Improving the performance of a complex environment will be an ever-present objective. As information processing becomes more accessible to the end user, so is it becoming more complex for the support staffs behind the scenes. BMC has made major inroads to help reduce that complexity.

ABOUT THE AUTHOR

Mary Lou Roberts is an independent consultant and free-lance writer. With more than 19 years of experience in the data processing industry, she has been involved extensively in both the technical and marketing aspects of the business.

Artificial Intelligence *from page 45*

Programming Language

One of the major limitations of shell systems is the lack of a powerful method for representing deductive thought processes. For this reason, it still continues to be desirable to use one of the powerful programming languages associated with artificial intelligence: PROLOG or LISP. These languages lack a standard user interface or a run-time facility. If, for example, rules are encoded in PROLOG, a user interface program must be developed to inspect the operation of the rules. As a result, languages such as PROLOG and LISP tend to be labor-intense and are less suitable to business application than are expert shells.

Programming Language Toolkit

Toolkits attempt to bridge the middle ground between expert system shells and programming languages. The facilities offered by a toolkit can vary across a wide range. They can range from extensive to austere. Based on the facilities offered, toolkits can look similar to expert system shells or they can be little more than a programming language. Ideally, toolkits are a complete developmental facility including a programming language, a flexible way of representing and storing factual and judgmental knowledge and an inference engine to solve problems easily. Such toolkits may utilize or work in conjunction with PROLOG or LISP, but the newer toolkits are more likely to run on a personal computer and use languages such as "C" or Pascal.

Expert System Shell

Origin

Many of the original expert system shells originated with expert systems that were devised to satisfy a specific application problem. The applications were stripped of knowledge and altered for general purpose use. The name "shell" originates from this ancestry: shells are applications that are devoid of any knowledge . . . an empty shell. However, many of the newer expert system shells are developed specifically as shells in an attempt to overcome some of the shortfalls of the original software.

The newer shells are typically either mainframe-based or personal computer-based with an interface into mainframe databases. The number of personal computer-based systems far outnumbers the mainframe-based systems. Further, the cost of personal computer-based systems ranges from $100 to $6,000 while mainframe-based systems range from $25,000 to $200,000. Mainframe-based systems are more robust; they support multiple users, more rules and larger knowledge bases.

Organization

The typical expert system shell consists of four components: a knowledge base, an inference engine, a user interface and a knowledge encoding facility. Shells are developed with languages such as LISP or PROLOG, but it is becoming increasingly more common for them to be developed in "C" or Pascal for personal computer based-systems and sometimes even COBOL for mainframe-based systems. Expert systems no longer require specialized hardware and it is now practical to use other languages.

Knowledge Base

The knowledge base is a collection of heuristic rules of behavior, the knowledge of good practices or good judgement that is common knowledge among the practitioners or experts in a field of knowledge. This collection of heuristic rules forms the base of knowledge and rules that experts use to solve problems. Sometimes an interactive editor is provided for developing the knowledge base.

Inference Engine

The inference engine is the computer software that replicates the logic or thought process of an expert. It is the compiler or interpreter that translates the base of rules in the knowledge base into executable rules. Accepting input from the expert system user and the knowledge base, the inference engine simulates the reasoning of the expert.

User Interface

The user interface is the computer software that allows the expert system user to enter facts and ask questions of the expert system. It is an executable system that allows the expert system user to interactively apply the rules to the knowledge base. The interface with the system is in a natural language.

Knowledge Encoding Facility

The knowledge encoding facility is the software used to acquire and encode the expert knowledge on the knowledge database.

Advantages

Expert system shells are the easiest way to prototype an expert system. The knowledge engineer does not have to contribute anything to developing the interface between the rule base and the knowledge base. Further, the knowledge engineer does not have to develop an executable system. This is especially important later since the executable system provides facilities that are common across most any expert system application. The execution time system facilities include such features as provisions to:

- Expand the meaning of a question
- Change the answer to a previous question or previous questions and propagate the impact through the session
- Volunteer information or answer questions before they are asked
- Inspect the knowledge base and to see previous or derived answers
- Justify why a question is asked or why an answer is given.

These common features are labor-intense to develop and their presence accelerates the prototyping and development of the expert system.

Disadvantages

The major limitation of expert shells is the limited power of their production rule system. For this reason, many expert systems continue to be developed using artificial intelligence programming languages such as LISP or PROLOG. These languages lack the features of the execution system components and are not as suitable for prototyping as an expert system shell. However, the search mechanism, data structures and interpretative nature of such languages as LISP and PROLOG do make them better tools than conventional languages such as Pascal and "C." Further, programming environments or toolkits are available to mitigate some of the negatives of LISP and PROLOG. When toolkits and shells are used in synergy, the result is the best of both worlds.

Selection Process

The process of selecting an expert system tool is a multiple step process. First, the organization must be aware that the implementation of expert system applications is possible and that the results are valuable. It requires a commitment on the part of the organization to use the results. It requires that the organization be prepared to make a long-range investment in one or more knowledge engineers, expert system software tools and, more importantly, an investment of from $2 million to $5 million in the development process.

The second step is the selection of the knowledge engineer. Knowledge engi-

CMG'89

Attend CMG'89

Bally's Hotel
December 11-15

Technology is changing faster than ever, creating unique challenges.

CMG'89 offers both technical and management sessions to help you keep pace with the latest industry developments.

Join more than 2,000 performance management professionals — you cannot afford to miss this information-packed conference!

RENO

The Premier Performance Management and Capacity Planning Conference

Industry leaders will address:

- **Network measurement**
- **Management reporting**
- **Capacity modeling**
- **DASD performance**
- **DB2/CICS/IMS/TSO**
- **Plus, much more**

Vendor presentations and display booths will showcase performance and capacity products.

For more information on this important conference, simply complete the form below and return to CMG headquarters.

Computer Measurement Group, Inc.
International Conference

I'd like more information on the conference to be held December 11-15, 1989.

Please send me a CMG '89 Preliminary Agenda.

Name: _____

Title: _____

Organization: _____

Street Address: _____

City: _____ State: _____ Zip: _____

Country: _____ Phone: _____

Clip and return to:

CMG '89
111 E. Wacker Dr.
Suite 600
Chicago, IL 60601
312/938-1228

neers are rare. Identifyng the correct skill set and selecting the correct person is difficult for the information technology group. In many cases, it is necessary to identify someone willing to learn the discipline and to make an investment in education. Both alternatives are risky, but the knowledge engineer is the keystone to the whole process of building an expert system.

The next step is selection of the expert system application. Since the technology is still quite limited, it is important that the correct application be selected. If an inappropriate application is selected and the process bogs down, the organization loses interest quickly. Further, the earlier in the process that a failure occurs, the greater the likelihood that the process will be abandoned. It is, therefore, critical that care be given to the application selection.

At this point, the knowledge engineer selects an expert system development tool. The knowledge engineer, you will remember, is an expert who specializes in understanding the strengths and weaknesses of the expert system tools. Based on the understanding and the technical environment, the knowledge engineer selects an expert system shell or toolkit to implement a prototype system and eventually the complete system. In some cases, the experience gained from the prototype causes a change in the tools used for the final system. Table 1 lists some of the common commercially available expert system shells and expert system toolkits.

Closing Comments

The skills and procedures required to develop an expert system differ considerably from those required to develop a commercial computer-based application system. The process of developing expert systems depends heavily on the skill of a knowledge engineer. The knowledge engineer performs the crucial steps of selecting the expert system application and the expert system tools.

Another crucial element is selecting a willing expert. The correct person is the expert who is indispensable to the organization. The knowledge engineer and the expert work together closely; therefore, the expert and the knowledge engineer need to establish rapport. Experts tend to embrace the expert system. They are a part of the entire process: the prototype development, system development and performance evaluation. However, convincing the non-expert to embrace the system is fraught with all the pitfalls associated with installing any new system. Here too, the key to success is to make the user part of the process.

In closing, it appears that the two most important aspects of building an expert system are the knowledge engineer and the choice of application. A knowledge engineer is a facilitator who understands the expert system development process and has the interpersonal skill to work with the expert, management and the user. The knowledge engineer works hand in glove with the expert, management and the expert system user to select the correct tools and the appropriate application.

FIGURE 1

Expert System Software

Expert System Toolkits

Goldworks
Gold Hill Computers
163 Harvard Street
Cambridge, MA 02139
617-492-2071

Texas Instruments Series
Texas Instruments Inc.
12501 Research Blvd.
P.O. Box 2909
Austin, TX 78769
800-527-3500

TOP-ONE/ai
Merrion, Inc.
World Trade Centre, Suite 400
Boston, MA 02210
617-439-5383

Expert System Shells — IBM PC Based

EXSYS
EXSYS, Inc.
P.O. Box 75158
Station 14
Albuquerque, NM 87194
505-836-6676

GURU
mbds, Inc.
P.O. Box 248
Lafayette, IN 47902
317-447-1122

VP-Expert
Paperback Software
International
2830 Ninth Street
Berkeley, CA 94710
415-644-2116

Copernicus and M1
Technowledge, Inc.
1850 Embarcadero Road
Palo Alto, CA 94303
415-424-0500

Levels PC
Information Builders
1250 Broadway
New York, NY 10001
212-736-4433

NEXPERT OBJECT
Neuron Data, Inc.
444 High Street
Palo Alto, CA 94301
415-321-4488

1st Class Fusion
Programs In Motion
286 Boston Post Road
Wayland, MA 01778
617-358-7722

PC Easy and PC Plus
Texas Instruments Inc.
12501 Research Blvd.
P.O. Box 2909
Austin, TX 78769
800-527-3500

Expert System Shells — IBM Mainframe Based

KBMS
AI CORP
100 Fifth Avenue
Waltham, MA 02254-9156
617-890-8400

Application Expert
Cullinet Software
400 Bluehile Drive
Westwood, MA 02090
617-329-7700

S.1
Technowledge, Inc.
1850 Embarcadero Road
Palo Alto, CA 94303
415-424-0500

AION DEVELOPMENT SYSTEM
AION Corporation
101 University Avenue
Palo Alto, CA 94301
415-328-9595

Expert System Environment,
Knowledgetool and APLPIE
IBM Corporation
(See your local IBM Marketing Representative)

ABOUT THE AUTHOR

Howard W. Miller, CDP, CSP, is responsible for administrative computing at Boston University. He has held senior-level positions in systems management for more than 20 years.

Effective DASD Growth Control

By Randall R. Lebedz

Lately, many articles have appeared in major technical publications proclaiming prolific growth in DASD networks of large mainframe shops. The *fact* that the growth has indeed occurred, I will not dispute (the bottom line of the data center manager's monthly hardware budget tells the true story). The *reasons* DASD growth is out of control will be the focus of this article.

How many times in recent months have you read that DASD requirements are growing at the rate of 20 to 35 percent per year? The interesting fact is not that the requirements for the storage of on-line (DASD) data are actually growing at this rate, but rather that the capacity of DASD on the floor in large mainframe shops is increasing by this rate. What does this indicate? It shows that DASD capacity is increasing at a much faster rate than actual DASD requirements. You might ask why is there such a notable difference between these two growth rates? The answer lies in the fact that the common solution to DASD space problems, as well as I/O subsystem performance problems, is to "throw more DASD at it." It is often too easy to give in to this temptation. You can easily justify (in your own mind anyhow) that since the DASD networks of most other large mainframe shops are growing at this rate, it is acceptable for yours to grow in this fashion as well.

Sooner or later, however, either the data center is going to run out of floor space or someone in upper corporate management is going to question the enormous expense of DASD in the operations budget. At that point you will be hard pressed to try to sell management on the old "cost-per-megabyte is cheaper" routine that *you* fell for. The person asking the question is looking at the bottom line and could not care less about cost-per-megabyte.

Hardware vendors will be the first to agree that storage requirements are escalating and that floor space is at a premium. They will also offer a readily available solution: replace existing hardware with more dense technology, thereby avoiding floor space problems. The drawback to this solution is that bottom line cost on these dense devices is not cheap and questions exist as to the unproven reliability of these new drives. Secondly, and even more important, is the question of whether or not you can believe that these new highly dense devices can be loaded up with data and achieve at least the same, if not better, response times as their predecessors. The vendors are claiming that these new, more dense devices can indeed achieve better response times, but their claims require a closer examination.

When migrating data from single or dual density devices to triple density DASD, it is of the utmost importance to evaluate the device-busy rates of each lower density source volume that will be combined onto the one triple density output volume. The source volumes' device-busys must be measured during the peak non-reschedulable period at your site.

When combining volumes, the device-busys are summed, not averaged. This means that if your three source volumes are each 10 percent busy, the new triple density volume you are creating will be approximately 30 percent busy. You cannot drive DASD, under normal circumstances, at rates above 15 to 20 percent busy without negatively impacting response times. The increased number of files in this case would most likely cause excessive queuing, excessive seeks, increased rotational delays and, therefore, increased disconnect time. This all translates to increased response times for users to the data in question. As previously indicated, this would be under normal circumstances.

Naturally the vendors can supply a fix for the aforementioned problems. It comes in the form of high-priced controllers offering four-pathing mode to the volumes, as well as higher-priced cache areas within those controllers. This is not to imply that there is no place for these controllers and cache areas in your storage subsystem, but rather that the cost-conscious storage management professional will employ a plan that will selectively utilize these expensive resources.

Bringing in DASD that has a lower cost-per-megabyte and then needing to install expensive controllers and cache to allow the new DASD to perform as well as the existing hardware seems counter productive.

A better alternative to using new hardware to deal with the rapid expansion of disk oriented data is to develop a storage management system that will allow customized control of data from a centralized area. The following account details the basic milestones which must be accomplished in order to curtail the rapid growth rate of DASD as well as provide the

expected rates of service as required by the users.

Establish A Storage Management Group

There is a great misconception in the DP world that storage management is a function that anyone can do. Many shops place the responsibility and the function of storage management in the computer operations area where it becomes an extension of the operations support group. The storage management systems that develop in these areas are never the type of systems which will hold the line on DASD increases. Storage management at this level consists of restoring datasets and resolving space abends.

Many installations have systems people doing "DASD management" as fill-in work around their regular jobs. There is no doubt that these are the sites where the DASD growth is at the 35 percent level. Storage management is a full-time job and unless management realizes this fact and dedicates the resources necessary to do the job correctly, the result will be a small savings in salary and a large long-term expenditure in hardware costs.

A storage management group must be established within the technical services department of the organization. This group must be acknowledged as being responsible for the storage and management of the company's data. This entity can consist of one person or of a number of people. The factors which will determine the number of people required to successfully manage the storage subsystem consist of (1) the software packages available for use at the installation, (2) the amount of DASD to be managed, (3) how elaborate a storage management system you want to develop and (4) how dynamic the senior storage management professional is at your data center.

Acquire Appropriate Software

A storage management software package that provides all of the key data management functions must be acquired. Any package can be used, even a "homegrown" one, but it should contain the ability to provide the following functions:

- Customized reporting of all datasets on DASD
- Support for all dataset organizations (PS, PO, DA, VSAM and even ISAM)
- Support for both ICF and non-ICF catalogs
- Support for all datasets regardless of catalog status
- The ability to migrate data from disk to disk
- The ability to archive data from disk to disk and/or tape/MSS
- The ability to archive/backup data in a compressed format
- A PDS compression function
- An idle space (dead space within a file) release function
- The ability to perform incremental backups of only changed data
- An auto restore function
- A flexible, easy-to-use control language
- A product that does not require the acquisition of additional DASD.

The number of storage management personnel required to manage the storage subsystem increases as the capacity of the storage network increases, but personnel requirements increase at a decreasing rate. This is due to the fact that once the framework of the basic storage management system is laid out, more volumes can be added to the system with minimal work by storage management personnel. The software will handle the increased number of DASD volumes.

Establish DASD Standards

DASD standards must be established, endorsed by management and incorporated into the *Data Center User's Guide*. These DASD standards must include efficient dataset naming standards. If you cannot readily identify the data, then you cannot easily manage it. Standards should be written to regulate the allocation of datasets in order to create a manageable environment.

The standards must prohibit volume ownership. This old concept creates management as well as performance problems. Data must be volser independent. Breaking the ties between datasets and specific volsers is the major stumbling block for installations attempting to go to a system managed storage type of environment.

Criteria for archival from disk to less expensive storage must also be addressed at this time. A realistic allowable period of inactivity on disk before archival must be established. For example, anywhere between seven and 35 days would be acceptable. Periods of less than seven days will tend to trigger excessive tape mounts for recalls to disk. If the threshold is set much over 35 days, few files will be archived after the original archive run. Archiving to disk in a compressed format is a possible alternative to archiving to tape or mass storage. However, archiving to disk should be limited to small datasets. You can expect to save about 50 percent by archiving to disk in a compressed format. On the other hand, you are still wasting a large amount of disk space on inactive data. I take the position that if you have not accessed a file in 35 days, then you can wait five more minutes for a tape mount. As far as the complaint that in some shops tape mounts take much longer than five minutes, this is simply a management problem in the operations area.

FIGURE 1

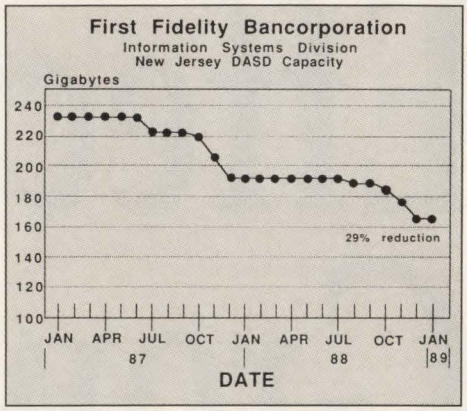

Define DASD Pools

DASD pools which logically group storage volumes must be created. For ease of implementation, DASD volumes should be created with meaningful volsers. Initially, pools can be categorized by function; that is, PROD Pool, Test Pool, TSO Pool and so on. The more sophisticated storage management administrator will not only define pools based on function as above, but also will incorporate performance requirements within each functional group. There are some effective pooling packages available to assist you in implementing the pooling concept: Empact Software's (Conyers, GA) Pool-DASD and Sterling Software's (Rancho Cordova, CA) VAM.

Create A Chargeback System

Reports must be created which will identify to management the applications' cost for use of storage resources. Establishing dollar values for different storage mediums will create a desire by users to utilize less expensive resources and avoid the anticipated pain that often accompanies device conversions. Implementing a chargeback system will provide an incen-

See DASD *page 85*

Seven Great Reasons To Join The BYTE Book Club...

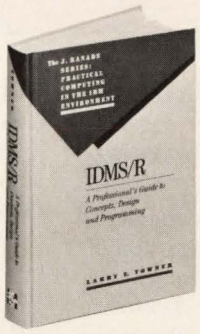

650/87X
Pub. Price $39.95

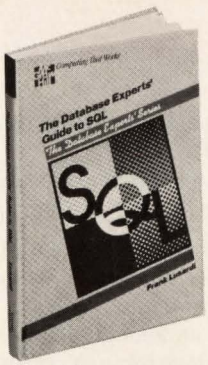

390/061
Pub. Price $39.95

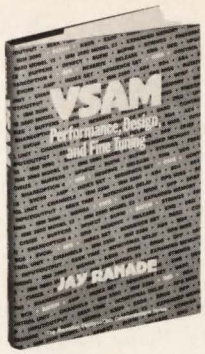

583963-6
Pub. Price $39.95

Take ANY ONE Of These "Large Systems" Books For $2.89!

(Values up to $114.50)

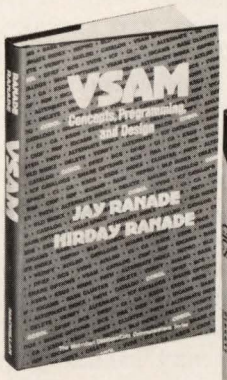

511/985
Pub. Price $39.95

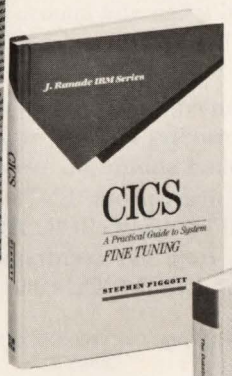

500/541
Pub. Price $39.95

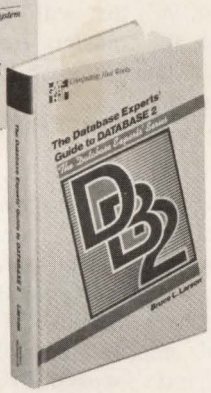

364/885
Pub. Price $39.95

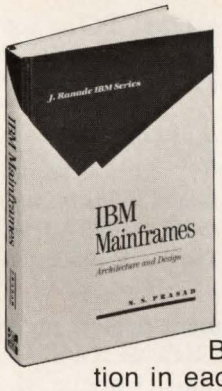

506/868
Pub. Price $39.95

And Get This $75.00 Encyclopedia ABSOLUTELY FREE!

- 1,047 pages
- $75.00 Publishers Price

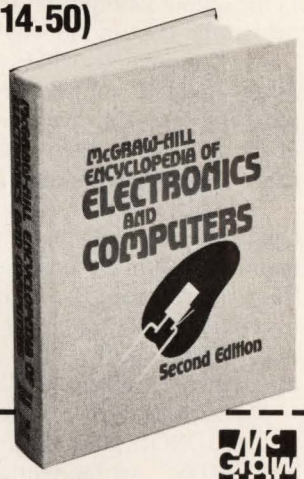

The BYTE Book Club introduces BYTE Large Systems . . . a special section in each monthly Club bulletin offering the mainframe and large systems books you need. The Club will still bring you the latest books on microcomputing, software engineering, languages, and peripherals, but now we also bring you the most up-to-date information on VM, CICS, VSAM, MVS, and more!

HOW THE CLUB WORKS TO SERVE YOU:

Quality books, dependable service.
You get a wide choice of books from *all* publishers, delivered right to your door. Every 3 or 4 weeks members receive the Club Bulletin, together with a dated Reply Card. If you want the Main Selection, do nothing—it will be shipped to you automatically. If you want another selection, or no book at all, just indicate your choice on the Reply Card and return it by the date specified. You'll always get at least 10 full days to decide. If the Club Bulletin ever comes late and you receive a Main Selection you don't want, *return it for credit at our expense.*

Substantial savings and Bonus Books. Discounts can run as high as 40% on every book you buy, and you're automatically eligible for our Bonus Book Plan, which allows you savings up to 70% off publishers' prices.

Easy terms. Your only obligation is to buy two more books during the next 12 months, after which you enjoy the benefits of membership with no further obligation. You or the Club may cancel membership anytime thereafter.

Write to: BYTE Book Club®, P.O. Box 582, Hightstown, NJ 08520-9959 or, for faster service, call **1-800-2-MCGRAW.**

McGRAW-HILL BOOK CLUBS
BYTE BOOK CLUB®
P.O. BOX 582 • HIGHTSTOWN, NJ 08520-9959

Please enroll me as a member and send me the book indicated, plus the McGraw-Hill Encyclopedia of Electronics and Computers 2/e, all for only $2.89, plus local tax, shipping and handling charges. I agree to purchase a minimum of two additional books during my first year of membership as outlined under the Club plan described in this ad. Membership in the Club is cancellable by me or the Club anytime after the two book purchase requirement has been fulfilled. A shipping and handling charge is added to all shipments.

Your **FREE** Encyclopedia: **454/99X**

Write Code No. of your selection here:

Signature _____

Name _____

Address / Apt. # _____

City / State / Zip _____

This order subject to acceptance by McGraw-Hill. All prices subject to change without notice. Offer good only to new members. Foreign member acceptance subject to special conditions. **P39831**

Q We are an MVS site running a Sales History to microfiche and have been missing data. The audit reports are correct and the microfiche vendor has rerun the job and is getting the same results — what is causing the problem?

A The importance of microfiche data in an MVS environment cannot be underestimated as an alternative to keeping data stored on disk resulting in heavy resource requirements for infrequently accessed data. Fiche data is also an alternative to tape datasets, eliminating the need to request data through a batch job when the user requests the data. When setting up a file to be microfiched, you must remember how the MVS environment will handle the process.

This particular process originated in applications six months prior to the problem being discovered by the user. The test results showed the same number of records being written out to the vendor's fiche tapes as were read in from the Sales History file. Spot checking was used to verify the microfiche output due to the size and amount of data involved.

In researching this problem, concentrate on three main areas: JCL for job creating fiche tape, TMS showing fiche tape as going off-site and the COBOL program used to format data for fiching. The JCL should reflect the DCB information as supplied by the microfiche vendor. TMS should reflect the number of reels being sent to the vendor as stated and in the same order as the output listing from the JCL. The COBOL program should format the data in the way it will be viewed on the fiche.

Also, remember what overrides were in the MVS environment. If DCB information is listed for a file in the COBOL program, this will override any reference to DCB information in the JCL and, thus, TMS will reflect incorrect DCB information as it gets its input from the JCL.

This problem was caused by the COBOL program that was copied from an existing program with DCB information larger than what the microfiche vendor could handle on the equipment. The JCL reflected the correct DCB information but was being overriden by the program. TMS was showing incorrect DCB information due to what the tape had cataloged, which was different from what the JCL had for the DCB.

The problem was resolved by changing the DCB information in the COBOL program to read ''BLOCK CONTAINS ZERO RECORDS'' for the output fiche file. This let the DCB information in the JCL determine how the fiche tape was going to be formatted and made TMS reflect the correct DCB information even though it was being sent off-site and not cataloged in the system.

Final thoughts on microfiche data should include the size of the data blocks being sent to the vendor. Many microfiche vendors cannot handle blocksizes over 10,000 bytes. Larger blocksizes cause truncation of data. Also, remember that if a COBOL program has DCB information coded for a file, it will override DCB information in the batch JCL and TMS will show incorrect DCB information for the file.

(Answer provided by Daniel A. Harris of Davis, Thomas & Associates, Minneapolis, MN)

Q I would like some information on AUTOSTATUS and when it should be used.

A AUTOSTATUS is a facility used with Cullinet's ADS/Online. It allows ADS/Online to return status codes to an issuing dialog. AUTOSTATUS is designated on an individual dialog basis in the Dialog Definition screen when the dialog is generated (ADSG). When AUTOSTATUS is designated, ADS/Online will only return certain common status codes (refer to table 10-1 *ADS/Online Reference Manual*) to the dialog. Only the status codes in table 10-1 will be returned; all other status codes will abort the dialog, unless AUTOSTATUS is overriden in the command line.

For the cases when the programmer would rather handle the error situation in the process code, AUTOSTATUS can be overriden by using the ''ALLOW'' clause in the command line. The ''ALLOW'' clause gives the programmer the capability to allow only specific expected status codes and will abort on those not specified. The ''ALLOW'' clause, however, can only be used when AUTOSTATUS is designated and handles only the codes included in the clause.

Since AUTOSTATUS can be overriden so easily, it can be used in most dialogs and in several shops it is considered a standard. By designating AUTOSTATUS for the dialog, the benefits are there if and when you choose to use them.

(Answer provided by Bob McDermott of Davis, Thomas & Associates, Minneapolis, MN)

Q I am planning on connecting my PC TOKEN RING LAN to my mainframe running VSE/SP and CICS. What are some of the considerations I should be aware of?

A In CICS, be sure that you put 'EXTDS' as one of the features for all PCs that will be using CICS. Also, if you are planning to use file transfers from the mainframe to the PC, in VTAM be sure that you use a LOGMODE table entry that has the query bit set on:

```
XXXXXXXX MODENT ...PSERVIC=X'XX80XXXXXXXXXXXXXXXXXXXX'
                              Query Bit
```

You should look into whether to use a gateway PC as your entry point into the mainframe. Each PC can act as a PU or, through use of a gateway, each can be an LU. There are advantages and disadvantages to both.

Be sure to call your IBM SE and request the INFO/SYS items with the keywords '3174 TRN PLANNING'. Also, I recommend ordering the IBM World Trade Manual, *Installation Guidelines for IBM Token-Ring Network Products GG24-3291*.

(Answer provided by Jerry Peterson of Davis, Thomas & Associates, Minneapolis, MN)

Q I am currently involved in a development project on an IBM MVS/XA 3090 using CICS 1.7, DB2 Release 3 and PL/1. When using CEDF to debug transactions, I notice DB2 calls are interrupted with the command ''call to resource manager.'' The address of the parameter list for the call is usually above the 16MB line. However, when I access most PL/1 dynamic storage it appears to be acquired ''below the line.'' The addresses for the DB2 storage takes four bytes and usually the format '80xxxxxx' (80 is a valid part of the address). The storage below the line commonly has the format 'NNxxxxxx' where NN appears to be FE, FD, FC,)) and so on. (Fx is not part of the address.) My question: How do you know when to use 24 bits and when to use all 31 bits, especially in CICS?

A First establish whether you want to run COBOL programs above or below the line. One reason to run above the line is to reduce the size of your DSA. COBOL II is also needed to run above the line. When you are linkediting a program to run above the line, you can specify RMODE=24, ADMODE=31. In answer to your question, check the mode setting of the programs' linkedit output.

(Answer provided by Dennis Bertrand, Davis, Thomas & Associates, Minneapolis, MN)

DASD *from page 82*

tive for application managers to reduce extra space in their systems.

The development and implementation of the above plan has allowed First Fidelity Bank (North Brunswick, NJ) to reduce its total DASD capacity by 29 percent over a two-year period. Monthly lease expenditures were reduced by an amazing 35 percent over the same two-year period (see Figure 1). During the same period, due to consolidations and increased business activity, the CPU workload grew more than 25 percent. This fact makes the disk reduction even more amazing! By establishing a corporate storage management methodology, getting the backing of senior management and applying the methodology in this article, it will be possible for you to produce similar results at your installation.

ABOUT THE AUTHOR

Randall R. Lebedz is a data processing officer in the Technical Services Department of First Fidelity Bank in North Brunswick, NJ. Lebedz is responsible for storage management and I/O configuration performance.

VSE/SP *from page 75*

For VSE systems running under VM, I use STDLABx for STDLABEL, STDLABSx for STDLABUS and STDLABPx for STDLABUP where the "x" at the end of the procedure name is the number of the VSE guest. This helps relate it to the IPL procedure ($IPLx) and ASI procedures ($$0JCLx, $$1JCLx and so on).

ABOUT THE AUTHOR

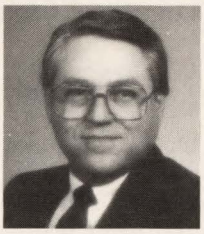

Mark Hanna is an IBM Business Partner. His consulting firm, Hanna & Associates, specializes in CICS, VSAM, COBOL, VM and VSE installation and support. He is the author of CICS Concepts & Facilities *and co-author of* Introduction To VSE & VSE Job Control *written with Suzan Hanna. Hanna has been in data processing since 1967. Hanna & Associates, P.O. Box 3325, Edmond, OK 73083-3325, (405) 340-1457.*

MVS *from page 70*

Not only did we manage to improve a visible part of the system's performance, we also gained a great deal of credibility and trust from our user community — a worthwhile objective at any time.

ABOUT THE AUTHOR

Patrick L. Gaul is a systems programmer responsible for performance analysis and capacity planning at Canadian Airlines International Ltd. He has been involved in the performance analysis and capacity planning fields since 1980, the last four years of which have been with Canadian Airlines.

Radik A. Gens is responsible for operating systems performance and capacity planning at Canadian Airlines International Ltd. His experience with computers dates back to 1964 when, as a student at a special mathematics school in the USSR, he began programming Soviet-made, zero-generation computers.

Canadian Airlines, Lansdowne Plaza, 4940 No. 3 Road, Richmond, B.C. V6X 3A5, (604) 270-5284.

ISPF 2.3

New features provide productivity improvements

By Jon E. Pearkins

Maximizing programmer productivity with ISPF continues to be the theme of this third in a series of articles in *MAINFRAME JOURNAL*. It is preceded by "ISPF Spells Productivity" in the November/December 1988 issue and "ISPF Techniques" in the February 1989 issue. Thanks to an unprecedented level of response from readers by letter, telephone and Reader Feedback cards both to the publisher and directly to me, there is enough new material for more articles.

In the first article, ISPF Version 2 Release 3 was described, sight unseen, from the Summary of Changes in the beginning of each ISPF 2.3 manual. Given six months of experience in actual use, as well as readers' experiences, this article encompasses only those new features of 2.3 that can provide productivity improvements.

Even More Productive

As well as adding the F and = Line Commands to The Most Productive Panel (see Figure 1), ISPF now *remembers* the DSNAME LEVEL and VOLUME you entered the last time you used it.

The change in wording from DISPLAY FORMAT OPTION to INITIAL DISPLAY VIEW is also significant. When viewing the dataset list produced by 3.4, you can use the LEFT and RIGHT PF keys (by default, PF10 and PF11) to switch between volume, space and attribute information. As well as saving key strokes over exiting, typing the desired DISPLAY VIEW and hitting ENTER again, there are major resource and response time savings gained by not repeating the look-up of each dataset in the catalog and VTOC again.

Almost all of my time in ISPF is spent in 3.4. Another subtle improvement in ISPF Version 2.3 eliminates one of the few remaining reasons to leave 3.4: using Edit to create the first member of an ISPF Library or PDS. You can now type the E Line Command, cursor over to the end of the dataset name on the same line and type the parenthesized name of the member you want to create. Alternatively, type E /(member) right over the top of the dataset name where "member" is the name of the member you want to create. As was mentioned in the first article, the slash ("/") tells ISPF to substitute the dataset name at this point in the command. VM users will recognize this as another good idea stolen from VM/CMS.

Undocumented Improvements

Sequence numbers, so valuable for change control, have also been improved. For a member with ISPF statistics but without valid sequence numbers, the Edit Primary Command RENUM is used to set the number of modifications in the member's statistics (MOD) to the number of lines (that is, all lines were considered changed by RENUM's addition of sequence numbers to each line). Version 2.3 has fixed that. RENUMbering is no longer considered a modification.

Two new Line Commands have been added to the ISPF Editor: Upper Case (UC) and Lower Case (LC). UC converts letters to upper case. Because it is not a single letter Line Command, UCUC or UCC must be used to mark the first and last line of a block of lines. LC is the opposite of UC, converting all letters in one or more lines to lower case.

UC and LC would have come in handy in earlier versions of ISPF. UC would have gotten you out of those situations where you typed a short document for printing, only to discover that the available printers could not print LC and were not set up with *case folding*: automatic translation from LC to UC. On the other hand, LC could have saved you a lot of retyping in those situations in which the current Edit Profile had CAPS ON, especially if you filled the screen before hitting ENTER. With LC, you would only have to overtype the first letter of each sentence and a few proper nouns and acronyms rather than overtyping 99 percent of the screen.

When determining what to do with your

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:
   INITIAL DISPLAY VIEW  ===> VOLUME       (VOLUME,SPACE,ATTRIB,TOTAL)
   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      F  – Free unused space
E  – Edit dataset         U  – Uncatalog dataset    =  – Repeat last command
D  – Delete dataset       P  – Print dataset
R  – Rename dataset       X  – Print index listing
I  – Dataset information  M  – Display member list
S  – Information (short)  Z  – Compress dataset     TSO command or CLIST
```

FIGURE 2

```
                    ISPF 3.14 — Search
                  SEARCH — FOR UTILITY
COMMAND ===>
SEARCH STRING
  ===>
MULTIPLE STRINGS ===> yes  (Yes to specify additional search strings)
ISPF LIBRARY:
  PROJECT      ===> YOURID
  GROUP        ===> ACTIVE   ===>          ===>          ===>
  TYPE         ===> CNTL
  MEMBER       ===>                       (Blank or pattern for member selection list,
                                           '*' for all members)
OTHER PARTITIONED OR SEQUENTIAL DATASET:
  DATASET NAME          ===> 'SYS1.PROCLIB(*)'
  VOLUME SERIAL         ===>               (If not cataloged)
DATASET PASSWORD        ===>               (If password protected)
LISTING DSNAME          ===> SRCHFOR.LIST
```

ISPF List or Log Dataset, you now have a fourth choice, KN. As you may recall, K means Keep the dataset and reuse it in the next session. KN lets you Keep the dataset, but allocate a New dataset in the next session. This allows you to have several datasets, one for each ISPF session that produced it. You might want to use KN when an ISPF session goes bad, but more pressing matters require your immediate attention. Typically, you would end the ISPF session and start another but keep the log and list datasets for later review.

Two Problems

Tom Rusnak, systems programmer at C.P.S. Direct Marketing in Phoenix, AZ, was one of several readers to phone or write, mentioning that uncataloging a tape or Generation Data Group (GDG) base entry from 3.4 (The Most Productive Panel shown in Figure 1) is no longer permitted in 2.3. In fact, there is even a new error message that will greet you when you type U for a tape or GDG: GDG Base or Tape Entry.

Henry Nalven of Marriott Corp. in Washington, D.C. phoned to say the problem was reported to IBM, but at press time, he had not yet received a formal response.

Rusnak, on the other hand, took a different approach. Because a new feature of ISPF 2.3 is to pass anything typed on the 3.4 List Panel that is not a valid Line Command to TSO for processing, he wrote two CLISTS to solve the problem. Members UT and UG were added to the PDS defined in the //SYSPROC DD statement found in the Profile used at TSO logon. UT Uncatalogs Tapes, while UG uncatalogs GDG base entries. Here is the code for them:

```
UT:
  PROC 1 &DATA
  DEL &DATA NOSCR
```

```
UG:
  PROC 1 &DATA
  DEL &DATA GDG
```

It is important to note that UT and UG are TSO CLISTs, not ISPF Edit Macros.

Searching A Library For A String

Finding all occurrences of a character string in all members of a PDS or ISPF Library had previously required the purchase of a non-IBM utility to enhance ISPF. ISPF Version 2.3 added option 3.14 (see Figure 2) as a new feature for just this purpose. Just type an asterisk (''*'') in the MEMBER field to indicate you want to search all members or leave it blank for a member list from which you can use the S Line Command to select the members that should be scanned.

When you type YES in the MULTIPLE STRINGS field, a panel (see Figure 3) is displayed permitting you to search for lines that contain two or more search strings.

FIGURE 3

```
ISPF 3.14 — Multiple Search Strings
SEARCH — SYS1.PROCLIB(*)
COMMAND = = =>
Specify 1 or more SEARCH STRINGS below:
  = = => _
  = = =>
  = = =>
  = = =>
  = = =>
  = = =>
  = = =>
  = = =>
  = = =>
  = = =>
Press ENTER to start search or END command to exit.
```

For a line to qualify as ''found,'' all of the strings must appear on that one line, but the order in which they appear is irrelevant. For example, if you are converting NCCF CLISTs to NetView, you would want to search for any lines that contain *both* the &WAIT keyword and the single quote (''''), no matter where they appear on the line. The Search Strings panel would look as follows:

```
= = => &WAIT WORD
= = => '''' C
```

WORD means that &WAIT must be delimited by blanks or punctuation and C means that the second line is a continuation of the search being specified on the first line. PREFIX and SUFFIX can also be used to indicate that the string will only be considered *found* if it is at the start or end of a *word*. If the search string contains single or double quotes or blanks, it should be enclosed in single quotes. Avoid the use of double quotes because experience has shown that searching for a single quote (as above) using double quotes as delimiters still requires a doubling of the imbedded single quote: ''''' is correct; '''' gives you an error message.

The listing you see as a result of your search (you are automatically put in ISPF Browse to view it), is put in a dataset. This means you can print it or browse it using ISPF's normal facilities at your leisure. Another search by you will overwrite the listing unless you enter a different name in the LISTING DSNAME field of the 3.14 panel.

Storage Management

The addition of the F Line Command to The Most Productive Panel (Figure 1) was mentioned in passing when I reviewed the new features of ISPF 2.3 in the first article. However, used diligently by ISPF users, it can have a positive impact on DASD space.

"Everyone hates getting B37 abends when creating datasets, so almost all datasets are over-allocated when created. This would not be a problem if all over-allocated datasets were freed by the creator after the dataset was loaded. This can be easily done . . . by entering an F next to all over-allocated datasets. This will free all unused space and leave the dataset using 100 percent of allocated space. This can be done to partitioned datasets as well as sequential datasets. If everyone used this technique, there would be considerably more space available," says David Levine, data administration analyst for Sony Corporation of America, Park Ridge, NJ.

Another area of storage management, for which the new capabilities of The Most Productive Panel can really help, is dataset migration. Systems such as IBM's DFHSM, Sterling's (Rancho Cordova, CA) DMS/OS, Computer Associates' (Garden City, NY) ASM2 and Innovation's (Little Falls, NJ) ABR provide a transparent RECALL (to use HSM terminology) of datasets that have not been used recently and migrated to less costly storage. *Transparent* is important. Unlike traditional archival methods, migrated datasets are automatically recalled from tape or compressed DASD, *without* endangering the successful execution of production jobs.

The ability to enter TSO commands beside datasets in 3.4 makes it much more convenient for users to migrate datasets they are unlikely to need in the foreseeable future. For those storage management systems that provide TSO commands for manual, user-initiated migration and recall, the user only needs to enter a command like HMIGRATE to the left of the first dataset on the screen that should be migrated and an equal sign ("=") beside any others.

Next Time

That wraps up the review of new and productive ways to use ISPF 2.3. In forthcoming articles, I will cover capabilities that have existed for some time in ISPF that readers indicate really help their productivity.

The topic drawing the most attention from readers, virtually untouched in the first two articles, is text formatting. As a result, the next article, "ISPF and Text," is dedicated to the topic of text with tricks for global search and global replace, as well as an in-depth look at text formatting.

ABOUT THE AUTHOR

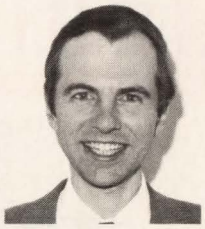

Jon E. Pearkins, CSP, CDP, is President of Certified Software Specialists Ltd., a Canadian company with a worldwide customer base for its line of IBM mainframe systems software. His 15 years of data processing include a variety of experience, both doing and managing in applications development, classroom and computer-based training, help desk, technical writing and data communications. Certified Software, 54015 Range Road 212, Ardrossan, Alberta, Canada, T0B 0E0, (403) 998-0607.

Don't take chances with systems education.

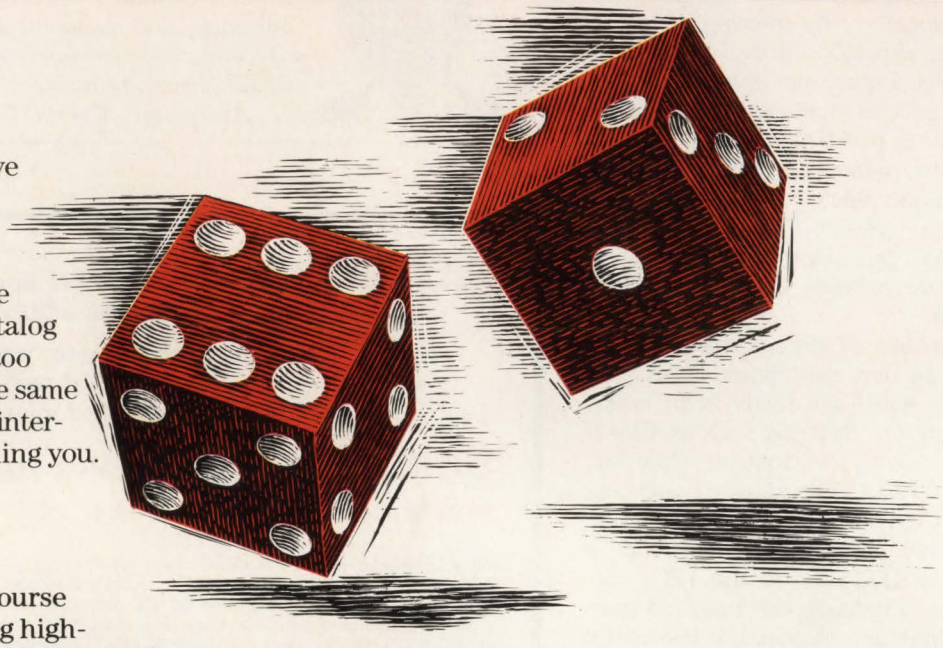

You've paid your money. You've taken time away from work. You're sitting in the classroom. And now you're stuck with "Lady Luck." Odds are the course doesn't live up to its catalog description. Or there are just too many students crowded into the same room. Or the instructor is more interested in selling you than teaching you.

Don't waste your precious time. Take Amdahl and be sure.

When you take an Amdahl course you can be 100% sure of getting high-caliber, professional courses, taught in well-equipped, comfortable classrooms by well-informed, objective instructors.

Amdahl's education service caters to all S/370 architectures, regardless of vendor hardware. Our classes are small, open forums where individual questions are easily addressed. Our staff is exceptional. Amdahl instructors have the technical expertise of veteran programmers and are effective communicators who teach complicated material clearly and comprehensively.

Don't settle for less. Take courses in fully-equipped Education Centers, nationwide.

All Amdahl courses are taught in full-service Education Centers located across the country. Most also feature on-line labs for hands-on application. We schedule over 500 classes in 50 different subjects each year at Centers in Chicago; New York; Columbia, MD; Atlanta; Houston; Anaheim and Santa Clara, CA; Washington DC; and Hartford, CT.

Don't take our word for it. See for yourself. Take an Amdahl course, risk-free.

If you want the best technical instruction in MVS, VM, IMS, DB2 or Communications Systems, your best bet is an Amdahl systems education course. We're so sure we'll live up to your highest expectations that we'll fully refund the cost of your Amdahl course if we don't.

Don't wait. Request a FREE 1989 Education Catalog, today.

For Express Enrollment information and a FREE 1989 Education Catalog, just give us a call. Or, complete and mail the postage-paid card attached. Don't wait. Classes fill-up early.

Don't take chances. Take Amdahl. Call today:

1-800-233-9521 ext. 135

Take Amdahl.

amdahl®
The SMART Choice

CIRCLE #60 on Reader Service Card ▲

Glaxo *from page 25*

building an Information Asset Directory (IAD) as part of its MIS Administrative Database. The IAD will contain all of the relationships between entities and the name of the generic relationship list transaction required to display the relationship. When this step is complete, the hard-coded relationships will be removed from the applications and replaced with table driven logic. This will allow driving the application from external data.

The other way these data driven design techniques speed up applications development is that the two basic types of applications can be cloned: entity maintenance and entity list. The basic logic remains constant; only validation edits and data name references need to be changed. Functional requirements can usually be handled by automating the navigation techniques.

For example, INSERT referential integrity checks leave the user with an ADD entity screen if a primary entity does not exist. An entity cannot be deleted until all mandatory relationships are removed and the application logic can move the user from one set of relationship lists to the next until all related rows have been deleted. When additional functionality is required, the standard cloned transactions can serve as a base for more specific functionally oriented transactions. So far, implementing transactions of this nature has not been necessary.

For the *decision support* applications, the rules are different. The real metric here is flexibility. Response times can be measured in minutes (although not *too* many minutes) and still be acceptable in many cases. However, *applications* to retrieve data must be quickly implementable without a traditional development cycle and tools must be available to provide ad hoc access to and analysis of existing data by lightly-trained (as opposed to totally untrained !!) end users.

For *batch jobs*, the performance criteria is pretty much the same as always. All batch jobs must execute within a fixed window of time, generally at night or on a weekend.

Even in a largely on-line environment, batch processing has a place. An excellent example of a function best done in batch is the posting of interface transactions. Glaxo has developed an Interface Management and Tracking (IMT) system allowing multiple applications on multiple platforms (across IBM, VAX and Hewlett-Packard hardware, using diverse file management software) to write interface information to a common collection structure.

The IMT structure itself is intelligent. At appropriate times it will activate the proper jobs to prepare interface data for cross-platform transfer and transfer it to the appropriate structures on the receiving platform, then kick off an IMT program on the receiving end after the transfer is complete. The receiving IMT software will post the transferred data to the appropriate places if it can. Then it will send the status of the transfer back to IMT software on the sending platform. The IMT is extremely modular, allowing it to be flexibly adapted to changing hardware and software requirements without undue difficulty.

Another area of high batch activity is the generation of summary files from detail data for end-user computing. And yet another area is traditional reporting. Many batch jobs of diverse types will be competing for the same finite time window. As much as possible, Glaxo's batch jobs are executable in parallel. Any job that must run alone is holding up other work that may be unrelated to it. Since some sequencing will always be required, CA-Scheduler is used to help manage the batch window with as little nightly operations staff interaction as possible.

Interesting Application: Image Scanning

This system is being addressed because it breaks new ground for Glaxo, both in functionality and in the amount of disk space that must be managed to support it. The application stores scanned graphic images as segmented Graphic Digitized Data Manager (GDDM) text in standard DB2 table rows. A traditional text-based application is integrated with the image storage/retrieval mechanism. This is all done using off-the-shelf hardware and standard software development tools.

First, some background information. The Glaxo Professional Relations department manages a Speakers' Program. This program is designed to provide speakers to talk about various pharmaceutical and medical topics. For instance, a sales representative will request a specific speaker or an available speaker for a particular date in a particular city on a particular topic.

Professional Relations personnel approve the request and coordinate the logistics of speaker selection and meeting arrangement. They also coordinate payment of expenses for meetings and speakers. In some cases a Curriculum Vitae for the speaker is sent to the sales representative prior to the assignment of a speaker. All of the information about the speakers is kept on paper.

The Recorded Information department had responsibility for storage of all the paper documentation. That department has a mandate to keep down physical storage requirements and to control the availability of information. This application was developed to help the Recorded Information department cut down floor-space requirements while improving availability.

During late 1987 Recorded Information Management requested a proposal on the non-paper alternatives available to manage the storage, indexing and retrieval of Speakers' Program documents. The volume was estimated to be between 300,000 and 500,000 documents. Glaxo looked at microfiche, optical systems and magnetic storage with DB2 in mind as the controlling DBMS.

It was found that image management technology was in its infancy for the IBM mainframe environment. After the preliminary evaluation of available alternatives, a short list of critical requirements which the various alternatives were tested against was drawn up. The requirements were the following:

- An image of each document had to be maintained on some media
- Each user of the system could use the same workstation to access documents and perform end-user computing functions such as word processing and spreadsheet work
- Response time for document retrieval had to be between 45 seconds and 1.5 minutes
- The documents would be indexed and accessible through any one or a combination of up to 15 keywords
- The application would share a CRT screen with an application that actually generated the documents
- The solution had to happen in six months and the cost could not be prohibitive (prohibitive was not quantified).

In order to develop an application that would meet these requirements, several technological issues were addressed and resolved. Number one on the list was software to manage document storage, indexing and retrieval. Images need to be digitized and compressed prior to storage.

TABLE 1

Partition Number	Begin Time-stamp	End Time-stamp

A PL/I program written elsewhere to convert scanned images to GDDM format was converted to COBOL and the file management routines were adapted to use DB2. The GDDM records are segmented into 4K pages before storage, then transparently reassembled when retrieved. 4K pages were chosen over the apparently more appropriate 32K pages under IBM's direction; apparently DB2 Version 1.3 has trouble managing 32K buffers. Had the system been developed under DB2 Version 2.1, perhaps 32K pages would have been used.

The DB2 tables are structured so that the images and index keywords are separated. Each image has approximately 27 rows of 2,000 bytes and an associated 1,000 bytes of index information. This brings up the major disadvantage to this approach: DASD requirements. Even after a compression algorithm is used, each image requires 56,000 bytes of storage, 55,000 for the image and 1,000 bytes for the index data. After a year of life and implementation of some archival procedures, DASD utilization for this application is expected to be between 15 and 20 gigabytes.

Second on the list is hardware to support the application. For scanning documents into the system, Glaxo used an IBM 3193 High Resolution Monitor and an IBM 3118 Scanner. The Monitor gives a crisp, clean image of the document as it is being scanned. The scanner, while slow, has been dependable. Scanning time depends on the size of the document with large documents taking proportionately less time per page than small ones. A one-page document takes 1.5 minutes to scan and index. A 35-page document (largest), takes 12 minutes.

For retrieving documents, the company uses the IBM Color Graphics Monitor, IBM 3270 Communication Board and Workstation Software. The communication board and workstation software provide for the presentation of the graphic image and splitting the screen. The screen is logically split horizontally. The image application is located on the bottom of the screen and a related traditional CSP/DB2 Speakers' Program application is located at the top of the screen. The user can hot key between the applications for inquiries and data entry.

There are a couple of interesting database administration aspects to managing large databases. Since Glaxo exceeds volume boundaries, partitioning is required, but partition management has proven to be non-trivial in this case.

The partitions are based on key ranges in a partitioning index. However, in order to simplify backup, the goal here is to completely fill one partition before starting to use another one and to be able to dynamically reassign the active partition as needed. Initially, 64 partitions were defined and the files were allocated for all but the first one at one track; the first partition was allocated at about 750MB. As the active partition fills, REORG the next one to its full allocation (roughly 750MB) and then direct the programs to make the resized one active.

Glaxo accomplished the dynamic partition reassignment capability by taking advantage of the fact that the image data records are automatically time-stamped with the current date and time. The partition number was defined as the clustering key in the data record, then an external partition mapping table was created as in Table 1.

Table 1 is read by the program storing image data and used to establish which partition number is moved into the partitioning key of the data record. The data record is then stored and will be written in the appropriate partition based on the contents of the partitioning key field. The partition mapping table can be maintained like any other DB2 table, thus to change partition assignment requires only adjusting the time-stamp range associated with a partition. No application code was written for this function; RC/Update from the Platinum Catalog Facility is used to make any changes.

Note that other indexes are defined to expedite data retrieval; the partitioning index is never used for retrieval, as it has no meaning. Note also that this approach is effective in this case because updates always happen through only *one* program; it might not be practical in other situations due to program maintenance overhead if something changes.

The other interesting point is backup. Glaxo does not have a large enough daily batch window to back up a tablespace the size of the image table using the automated mechanism described earlier, so this tablespace is in the NOBACKUP control table and is handled manually. The manual process consists of a daily full backup of the active partition and incremental backups of the other 63. Once a month the company bites the bullet and takes a full backup of the whole thing.

In retrospect, this solution to document management needs has been satisfactory. The microfiche solution would not have addressed the one workstation and interface requirements. The optical solution was and still is expensive and, at the time, was not a proven technology on the IBM mainframe.

But . . . Image management technology has progressed since the original proposal was presented. If Glaxo were to conduct a similar evaluation today, the solution might or might not be different, depending on the use of more canned software or optical storage instead of magnetic media.

Summary

This article described some of the reasons Glaxo is a successful DB2 shop. Hopefully, the information presented has been of interest, but the real basis of success goes beyond what is described here.

The company has succeeded as a DB2 shop for the same three reasons shops before Glaxo have succeeded without DB2. The first reason is vision. Glaxo knows where it wants to be and what it needs to accomplish to get there.

Management is the second reason. Management understands what it takes to get work done and provides strong and consistent support. Employees are provided with the tools and training needed to do the job and they are given proper credit for their accomplishments.

The last reason is competent, hard working employees. Vision and management are critical but insufficient. Good work gets done when you give a good idea to a good man or woman.

To summarize, you have to form a clear vision of what you want to do, convince management that it has a real payback on the bottom line, get a solid management commitment, then *go for it* using the best people you can find.

ABOUT THE AUTHOR

Robin Pasley is in charge of Database Administration at Glaxo Inc. His experience includes 10 years in data processing and five years in database administration. Glaxo Inc., 5 Moore Dr., RTP, NC 27709, (919) 248-2100.

COBOL *from page 36*

ger or the interactive debugger may be used. The batch mode debugger can be used in MVS, CICS or CMS. Before running a program with the batch mode debugger, a file of 80-byte records must be created that contains the DEBUG commands used during program execution.

The SSRANGE option checks subscripted areas, indexed areas and OCCURS DEPENDING ON areas to ensure that they do not exceed the storage size allotted to them by the compiler. Invalid values can cause storage overlays and/or protection exceptions (OC4s). For instance, if a one-dimensional table entry occurs 10 times and a program attempts to reference the 15th entry, an error message will be displayed. There is also an installation option that will terminate a program during execution if it gets this error.

Note that each subscript in a multi-subscripted table element is not checked for validity, only the resulting address. For instance, if a three dimensional table has a maximum subscript of (12, 10, 3) and you use a subscript that generates a value (1, 1, 6), no error will occur. Although one of the subscripts is outside its proper range, the address computed using all three subscripts will fall well inside the table limits.

While this option is useful, the machine code added by the compiler will make the program run considerably slower. Each time a subscript is used, even if it has not been changed, a COBOL subroutine is used to check it. While useful in development of programs, the SSRANGE option should be avoided in the production environment.

The RENT option makes a program reentrant. This allows one copy of a program to be shared by many users, such as is the case with CICS programs. The system implements this option by copying the areas of the COBOL program that are modified into a GETMAINed area. This includes the TGT and WORKING-STORAGE section. When the RENT option is used in conjunction with the DATA option, you may specify that program storage be acquired from unrestricted storage, either above or below the 16MB line.

PFDSGN is a new option that can make a program run more efficiently. It instructs the compiler that your numeric fields have valid signs and that the signs are F for unsigned fields and C or D for signed fields. Programmers often do not realize that A, C, E and F are positive signs while B and D are negative signs. None of these values will cause data exceptions (OC7s). However, when packed decimal calculations are performed on an IBM mainframe, the preferred signs of C and D are placed into the results. The NOPFDSGN option tells the compiler that numeric fields in a program do not necessarily have the preferred signs of F, C and D.

The FASTSRT option allows a program to process a sort faster. This is accomplished by having the sort run external to the program rather than having every record funnel through the COBOL I/O routines. It therefore only works when a SORT statement contains either a USING or a GIVING clause. If the SORT statement contains an input procedure and an output procedure, the FASTSRT option does not affect sort processing. In this case, FASTSRT does not cause an error if it is used as a compilation option, but you do receive an information message that the fast sort was not done.

Programmers should not overlook the fact that certain options will be required in a company's production environment. The options that affect program execution used in system testing should match the required options that will affect program execution in the production environment. This does not mean that programmers should be limited in their options used during testing. Companies should allow programmers to choose their own compiler options.

COBOL compiler options are meant to be just what their name implies, "options." Taking the time to understand the options is often neglected because knowledge of their use seems not to be required; a company usually has already prepared procedures to compile a program. However, the proficient programmer will tailor his compiles to meet his needs, thereby increasing productivity. Reduced costs should also be a favorable result of the additional knowledge.

ABOUT THE AUTHOR

Harvey Bookman is President of Bookman Consulting, Inc., a software development company specializing in programmer proficiency testing. Bookman Consulting, (212) 819-1955.

VM *from page 62*

totype or because of last-minute changes in the design that no one incorporated into the prototype. But most results should agree. This procedure helps guard against any psychological predisposition to accept unthinkingly what appears on the screen.

Batch reports may take more effort to verify because they tend to be longer and more complex. If the prototype produces identical reports, team members can COMPARE them on a minidisk. Normally, however, the prototype would not include all headings, footings and other formatting specifications. One way to deal with this problem is to use XEDIT's "search all" capability ("ALL /") to display and delete heading and footing lines so that the COMPARE will work. Another is to write an EXEC that parses each line to look for and compare valid data elements. Or, at a minumum, they can write an EXEC that searches through the reports for critical results. An EXEC will catch unexpected changes a human might not notice when glancing through the output.

The cost-effectiveness of writing such EXECs depends on the importance of accurate results. For simple jobs like internal mailing labels, it may suffice to look for the names of a few senior vice presidents. For customer account summaries, however, each failure of quality control can have a significant dollar price not only in good will and lost business, but also in the man-hours required to identify and correct errors. The EXECs will vary somewhat from report to report, but the project team's effort in writing them will be less than the time needed to correct the first production error.

This article is an excerpt from Chapter Ten of VM Applications Handbook *(McGraw-Hill, 1989) Gary McClain, editor.*

ABOUT THE AUTHOR

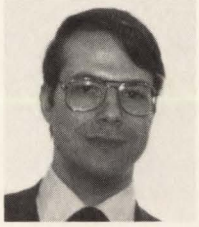

Michael Seadle, Ph.D., is head of user services at Eastern Michigan University, Ypsilanti, MI. He has worked as an applications developer, database administrator and VM systems programmer. Also, he helped establish and is an active member of the VM Enthusiasts of Michigan, a VM users' group. He may be reached at (313) 434-1576.

STROBE Now Attributes Resource Use In Service Routines

Release 8.0 of Programart's STROBE Performance Measurement System, an application tuning software product, now attributes resource use in system service routines to the user code that invoked each routine. This feature allows users to easily locate and fix areas of code that cause excessive overhead in computing resources. Other new features provide users with greater control over their measurement sessions, more intuitive methods for interacting with STROBE and more detailed information about measurement sessions. STROBE measures the performance of batch processing and on-line applications running in MVS/370, XA and ESA environments, including those that use CICS, DB2 and IMS and other vendors' subsystems.

For more information contact Martha Shaefer at Programart, Cambridge, MA, (617) 661-3020 or:

Circle #200 on the Reader Service Card

High-Performance Storage Subsystem Announced

Amdahl Corporation recently announced its new 6110 High-Performance Storage subsystem. The electronic unit, which employs semiconductor storage shareable among several host systems, is said to significantly improve response times, system throughput and productivity for users needing to access critical, highly active, on-line data. The 6110 can manage up to one gigabyte of data and the company believes its maximum aggregate data transfer rate of 36MB per second is the highest in the industry for this type of device. The 6110 can be attached to all S/370 compatible processors with data streaming channels, transferring data at rates of 3.0 or 4.5MB per second. It is available in four models. The 6110-10 and 6110-20, each with four storage adapters, respectively provide four and eight channels to host processors. The 6110-30 and 6110-40, with eight storage adapters, offer 12 and 16 channels respectively.

For more information contact Al Richard at Amdahl Corporation, Sunnyvale, CA, (408) 746-8829 or:

Circle #201 on the Reader Service Card

SCREENGEN Accelerates CICS Screen Development

SCREENGEN, a CICS screen development tool from MBA, Inc., eliminates the need to code CICS native IBM mapset definitions by providing a programmer-friendly "screen design" facility. Defined screens are then converted automatically into BMS mapset source code and associated COBOL copy groups. It operates with IBM, Panvalet or Librarian and handles all screen attributes including extended attributing. Screens can be easily modified (or cloned) individually or in mass and new mapsets generated. For detailed BMS code previously written, a companion utility program will convert existing mapset code into SCREENGEN formats for future use or modification.

For more information contact Carol Mersch at MBA, Inc., Tulsa, OK, (918) 587-1500 or:

Circle #202 on the Reader Service Card

DB2 ACTIVITY MONITOR Introduced

DB2 ACTIVITY MONITOR, from BMC Software, Inc., collects and displays both real-time and historical data for DB2 and all transaction environments. It generates batch reports from historical data and SMF as well as supplying a DB2 Console function. By locating inefficient use as it occurs, DB2 ACTIVITY MONITOR conserves CPU and I/O and minimizes CPU utilization by sampling control blocks and providing automatic trace control. It also increases system availability and reduces outages by isolating offending programs, identifying potential problems and allowing for quick resolution before they impact performance.

For more information contact Sandy Richardson, BMC Software, Inc., Sugar Land, TX, (713) 240-8800 or:

Circle #203 on the Reader Service Card

DB/REPORTER Now Provides Dual Support For DB2 & SQL/DS

Systems Center, Inc. recently announced the availability of DB/REPORTER Release 2.0, the first of its relational database products that now provides dual support for IBM's DB2 and SQL/DS database management systems. The dual support means that both MVS and VM can benefit from DB/REPORTER's report writing capabilities. Most notable of DB/REPORTER's new features is the ability to support flat file I/O processing that allows reports to include data from outside DB2 or SQL/DS and a much greater ability for producing reports and creating extract files for processing by other systems by reading and writing CMS files in the VM environment and sequential datasets in the MVS environment.

For more information contact Silas Matteson at Systems Center, Inc., Reston, VA, (703) 264-8000 or:

Circle #204 on the Reader Service Card

New VSAM Data Compression Package

Goal Systems' COMPRESSOR/VSE is a VSAM data compression package that can increase data density by using a sophisticated set of algorithms to identify and compress commonly occurring patterns in data. Depending on the application and type of data, savings of 10 to 90 percent in disk space are said to be possible. It supports KSDS, ESDS and VSAM-managed SAM. DASD and cost savings can be forecasted using the ANALYZER feature.

For more information contact Carrie Reber at Goal Systems International, Columbus, OH, (614) 888-1775 or:

Circle #205 on the Reader Service Card

ProAlter/Plus Provides Complete Set Of DB2 Tools

ProAlter/Plus, from On-Line Software, Inc., is a fully integrated solution for analyzing, maintaining and tuning DB2 systems. It allows DB2 database administrators to use ISPF-like panels to view and modify DB2 objects and optimize DB2 application performance. ProAlter/Plus also helps DBAs analyze DB2 catalog data in a hierarchical format through a series of "Show" and "Print" actions. Also, its Show Definition and PathAnalysis facilities assist DBAs in altering and fine tuning individual applications' access paths. Altering and migrating DB2 objects or sets of DB2 objects from test to production from one subsystem to another is said to be much easier with ProAlter/Plus.

For more information contact Steven Mariconda, On-Line Software International, Inc., Fort Lee, NJ, (201) 592-0009 or:

Circle #206 on the Reader Service Card

New Service Level Management Product

OMEGACENTER is Candle Corporation's integrated service level management system for data centers using MVS. Candle made OMEGACENTER possible by introducing Version 200 of the Status

Monitor and AF/Operator Version 200, with its tightly-coupled connection to OMEGAMON. The OMEGACENTER system is designed to make problem diagnosis and resolution fast and easy in order to maintain established service levels. The simplicity of the Status Monitor enables even less-experienced staffs to recognize service level problems at a glance.

For more information contact Kelley Murray at Candle Corporation, Los Angeles, CA, (213) 207-1400 or:

Circle #207 on the Reader Service Card

IMPACT Manages Changes In DB2 Production Environment

Infolink Software's IMPACT is a DB2 application product designed to enable corporate and data center management to manage changes and problems in a production environment. It also maintains the inventory and configuration for hardware, software and network components. IMPACT is a menu-driven system utilizing ISPF to provide a clear and concise list of options. These options guide users through the system, enabling them to select appropriate panels for controlling changes to be implemented.

For more information contact Ed Weiss at Infolink Software, Inc., San Mateo, CA, (415) 574-3305 or:

Circle #208 on the Reader Service Card

New DB2 Option For Optima's Change Man

Optima Software, Inc. has announced the availability of a new DB2 Option for its change implementation management system, Change Man. The DB2 Option is designed to enhance the productivity of programmers and change administrators responsible for making changes to DB2 applications. It will also guarantee the integrity of DBRMs and PLANs affected by these changes. In addition, Change Man's DB2 Option also provides valuable information about object relationships, audits the integrity of a proposed migration of changes from test to production, automates the migration process, dynamically inspects the DB2 catalog for inconsistencies between DBRMs and affected PLANs and prevents potential 818 errors by automatically reBINDing selected PLANs.

For more information contact Don Murphy at Optima Software, Inc., Sacramento, CA, (916) 646-3800, FAX (916) 646-3466 or:

Circle #209 on the Reader Service Card

CAT Scan: A Versatile LISTCAT Alternative

CAT Scan, a new product from Softworks, Inc., is an on-line LISTCAT alternative that produces an easy-to-read, condensed catalog listing of selected datasets. It eliminates manual searches through LISTCAT output and reports only on datasets meeting user-specified criteria, providing programmers with quick, efficient access to needed information. Users can choose from more than 70 dataset selection keywords to quickly find and report needed information. A user exit is provided for those who wish to access catalog information for more specialized analysis. It identifies datasets that are over allocated and datasets that have excessive CI/CA splits and/or extents. CAT Scan is compatible with IBM's MVS and Fujitsu MPS operating systems.

For more information contact Dave Krehbiel at Softworks, Inc., Clinton, MD, (301) 868-4221 or:

Circle #209 on the Reader Service Card

Shearson Lehman Hutton Uses Memory Management For I/O Relief

By Brandon McGowan

On-line transaction volumes at Shearson-Lehman-Hutton in New York often reach three million transactions per day across multiple IBM 3090-600E systems. One of the major problems associated with high volume transaction processing is erratic system performance.

The majority of Shearson's three million transactions use VSAM files for their data requests. Other transactions depend heavily on IDMS for their data. Both systems require heavy I/O activity which accounts for the bulk of the response time in the applications running on those systems.

Hardware vendors have attempted to resolve performance issues by increasing CPU memory sizes and cycle speeds. However, Shearson, like many users, has concluded that it is the DASD rotational speed that hampers applications. Adding more CPUs will not eliminate this problem. However, Shearson is taking advantage of recent advances in the size of memory and extended storage.

Shearson's solution has been to move those datasets which are accessed repetitively into virtual memory. Access to data in memory is one thousand times faster than access from DASD.

For the last few months, Shearson has been developing new systems using a different approach. Instead of keeping heavily-accessed data in a dataset on disk, the data can often be assembled into a table and accessed in memory. The software that automates this task is tableBASE from Data Kinetics (Ottawa, Ontario, Canada), and it has become a key resource for Shearson's data storage and retrievel applications.

Two applications being implemented at Shearson call for processing rates of up to 200 updates per second. These applications are primarily storing the most current market data rates for all securities being traded.

Richard Kneisz, Vice President of CICS Technical Services for Shearson, explains, "We can achieve this rate of throughput by letting tableBASE manage our data requests in memory. The response times we see on average are less than 0.0001 seconds per update."

The software, tableBASE, can work with any application that uses conventional IBM coding. All programs, whether written in COBOL, Assembler, a 4GL or some other language can share data that has been assembled into tables but with each having its own view of that data.

Kneisz adds, "Some applications we're developing would have required VSAM alternate indexes. Instead, we're using one of tableBASE's features called Alternate Views to achieve the same results with virtually no I/O cost, since most of the data is in memory."

Shearson has achieved substantial benefits in reducing batch run times by managing data in memory instead of on disk. Its applications often require data to be stored, summarized on several levels and put into numerous reports. Using ordinary sort routines required an extraordinary amount of file passing and I/O. The bulk of Shearson's I/O was eliminated by using a memory-resident sort facility. This means Shearson's programs only have to read data in once. Data is sorted during loading and summarized in memory without additional disk access and is made available to their report-generating routines.

Kneisz recalls, "One process was changed from using a VSAM KSDS file to using a tableBASE table and run time went from 2.5 hours to 10 minutes. The I/O count went from 300,000 to almost no I/O. The memory-resident sort facility was a major factor in eliminating an I/O bottleneck for us."

ABOUT THE AUTHOR

Brandon McGowan is a free-lance writer based in New York City.

Technical Support Services

DTA, the Midwest's largest Technical Support Consulting firm, has openings in two locations. DTA provides IBM customers with contract support services and system software:

Senior Systems Programmers
- *Openings in Minneapolis and Chicago*

- System Software Installation & Support (MVS, VM & VSE)
- DBA Installation & Support (IMS, DB2 & IDMS)
- Operations Consulting & Support (MVS, VM & VSE)
- Telecommunications Planning & Support (VTAM & NCP)
- Advanced CICS Technical Application
- Teleprocessing Monitor Support (CICS, IDMS-DC, IMS-DC)
- Capacity Planning, Problem Determination
- Remote Dial-up Support

Software Developer
- *Opening in Minneapolis*

Requirements of the candidate include:
- 5+ years of programming in BAL and some of the following skills:
- Basic understanding of CICS internals
- Basic understanding of VSAM
- MVS internals

Software Customer Support Rep
- *Opening in Minneapolis*

- VSE, CICS, VTAM Systems Programming (Background at least 1 to 2 years)
- Must enjoy assisting systems programmers

DTA offers a comprehensive salary and benefit package. If you desire to be the best in your technical field and assist numerous customers with implementing their technical solutions, contact either of our offices:

550 Waterford Park
505 North Highway 169
Minneapolis, Minnesota 55441
(612) 591-6153
Contact: Debra

103 Midwest Plaza South
2021 Midwest Road
Oak Brook, Illinois 60521
(312) 620-1322
Contact: Harold

DTA®

An Equal Opportunity Employer
No Agencies Please

Bringing Peace of Mind to Customers through Technical Support

Data Processing

Behind The Scenes— *Where The Real People Are!*

At Warner Bros., Inc., we believe you don't have to be the one in front of the camera to be recognized. We invite you to consider this challenging opportunity:

IBM SYSTEMS PROGRAMMER(S)

Minimum 2 years experience, MVS/XA, SMPE, VTAM, JES2, 370/ASSEMBLER. Data communications or PC experience a plus. Must be able to write code.

Home terminal supplied. Shirt sleeve environment, paid overtime, flexible hours.

For consideration, please send resume with salary history to:

WARNER BROS., INC.
4000 Warner Blvd.
Bldg. #30
Burbank, CA 91522

Equal Opportunity Employer M/F

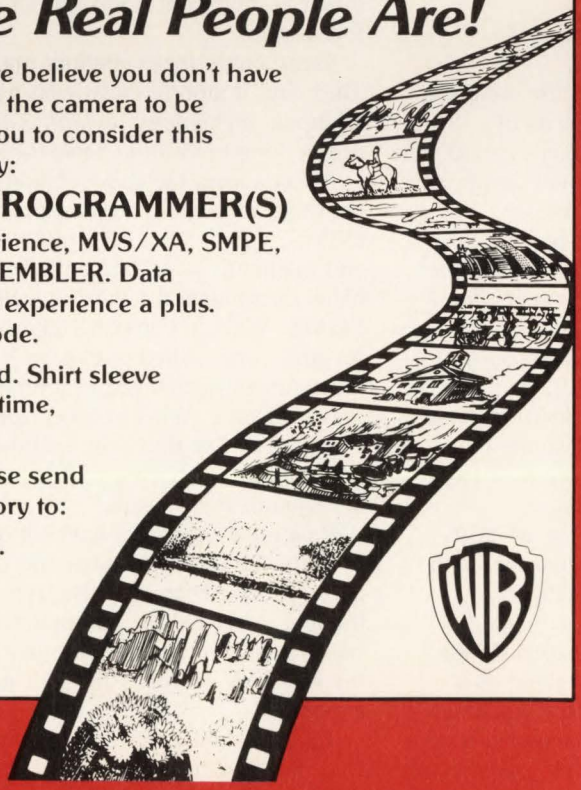

Systems Software Challenge

Software Pursuits is a leader in creating systems software solutions for the IBM mainframe market place. We are looking for bright, friendly people to help us create innovative system software of the highest quality. Since 1975 we have provided advanced products including our operating systems and other systems software products.

Product Development

We are looking for only the most talented professionals to work on some of the most complex and interesting problems in system software. Your programming skills must be instinctive. Give your imagination an opportunity to be expressed. Candidates with CICS, VTAM, Networking, 370 Assembler, or VSE mainframe experience are strongly preferred.

Software Pursuits offers an outstanding work environment, where your efforts will be appreciated. We promote a relaxed team spirit where achievement of the highest quality products and services is our primary goal. Please send resumé to:

SOFTWARE PURSUITS

Attention: Bill Boeckmann
1420 Harbor Bay Pkwy., #200
Alameda, CA 94501
(415) 769-4900

The Myths Of Relational Database

By Bill Backs

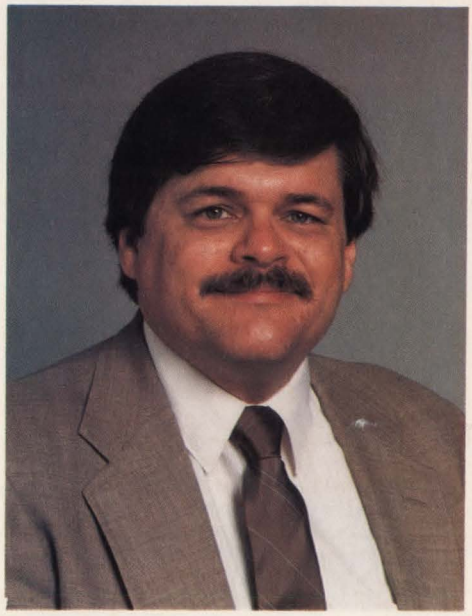

Bill Backs is Director of Information Technology at Scott, Foresman & Co. and President of the International DB2 Users Group (IDUG).

Two myths have grown up around relational databases. As with all myths, they contain a kernel of truth, but in general they have little to do with reality. Unfortunately that does not stop people from believing them and then acting on them as they perform their jobs. As a result, companies are failing in their application of relational technology to their businesses. In many cases they are then blaming the technology, instead of the myth.

Myth 1 — Relational is Easy

I remember attending a marketing presentation in the fall of 1985 given by IBM on the just announced database manager, DB2. The speaker told us that since DB2 was relational, it was so easy to implement and use that we would not need any DBAs. Imagine, having the latest in technology and being able to get rid of all those bothersome and expensive DBAs on your staff at the same time. I expect some people left that briefing actually believing it was possible. IBM quickly dropped that sales line, but the myth has continued. Put in its most basic form it is this: *relational databases are so easy to use, intelligent and flexible that it is no longer necessary to apply the rigorous technical skills and methodologies we as an industry have developed over the past 20 years. The software will handle it all.*

This myth is most frequently applied to the issue of database design. The process of database design, normalization and performance tuning is really quite complex, particularly when hundreds of entities and thousands of attributes are involved. It is as important for the analysts to know their data and how it is used within relational design as it is within any database design methodology. Entity-relationship modeling is a rigorous and exacting process. The results of the process, if it is done well, are efficient, flexible database designs which will serve a corporation's needs for years. Yet many people believe that because relational databases are so flexible (after all, they do not have any physical pointers, no PDBs — they are relational), they can ignore design or go back and do it ''later.'' Corporations that think nothing of spending six months or more designing an IMS database expect the process to take only a week or two under DB2. Then when they experience performance problems, when they cannot get to the data they want easily, when the database requires constant redesign, when the myth trips them up, they use their experience to show that relational technology is ''just not ready yet.''

The myth carries over from database to application design. In other databases the programmer has to worry about data navigation, that is, how to move around within the database to most efficiently access the data needed for the application. Now the database software does all that for the programmer. To a large extent that is true. But programmers who completely ignore data access issues are asking for trouble. They almost always find it. Then they also say, ''Well, it is DB2, what kind of performance can you expect?'' They have failed to separate the myth from reality.

Myth 2 — All Databases Are Alike

Several vendors are marketing ''bridges'' which allow relational databases to be accessed by programs originally written for other data storage methods. They allow, for example, programs written for VSAM file structures to read and update DB2 tables — but only if the DB2 tables are identical to the VSAM files in design and content. Application software vendors use I/O subroutines in their products so that only the data access routines have to be changed to port their software from one database to another. Once they have an IMS version of their product, they can churn out a DB2 or an IDMS version in a few weeks.

Surprisingly these vendors find a market for their product. They find it among those who believe the last myth of relational technology: *actually, all databases are alike; they are nothing more than sophisticated access methods.*

As an access method, relational databases are terrible (so are other databases, for that matter). When you purchase a DBMS, you acquire not only software which physically stores and maintains your data on disk, but also which manages and helps manipulate the data as well. That is the ''M'' in DBMS. When a company takes data structures it has created in some form, whether that be VSAM, IMS or whatever and moves them unchanged to a relational database, it has made a grave mistake. When it takes applications which were designed to maximize the facilities of a given access method and uses them unchanged against that new database, it compounds that mistake.

Relational databases have unique design considerations, both for the data structures and for the applications which access them. They must be taken into account and utilized or the end result will be an application poorly designed and which performs poorly. But then, who worries about design in the relational world? And isn't poor performance to be expected?

Take the time to do the job right. This has always been true. Relational technology has not changed it.

Working with DB2?

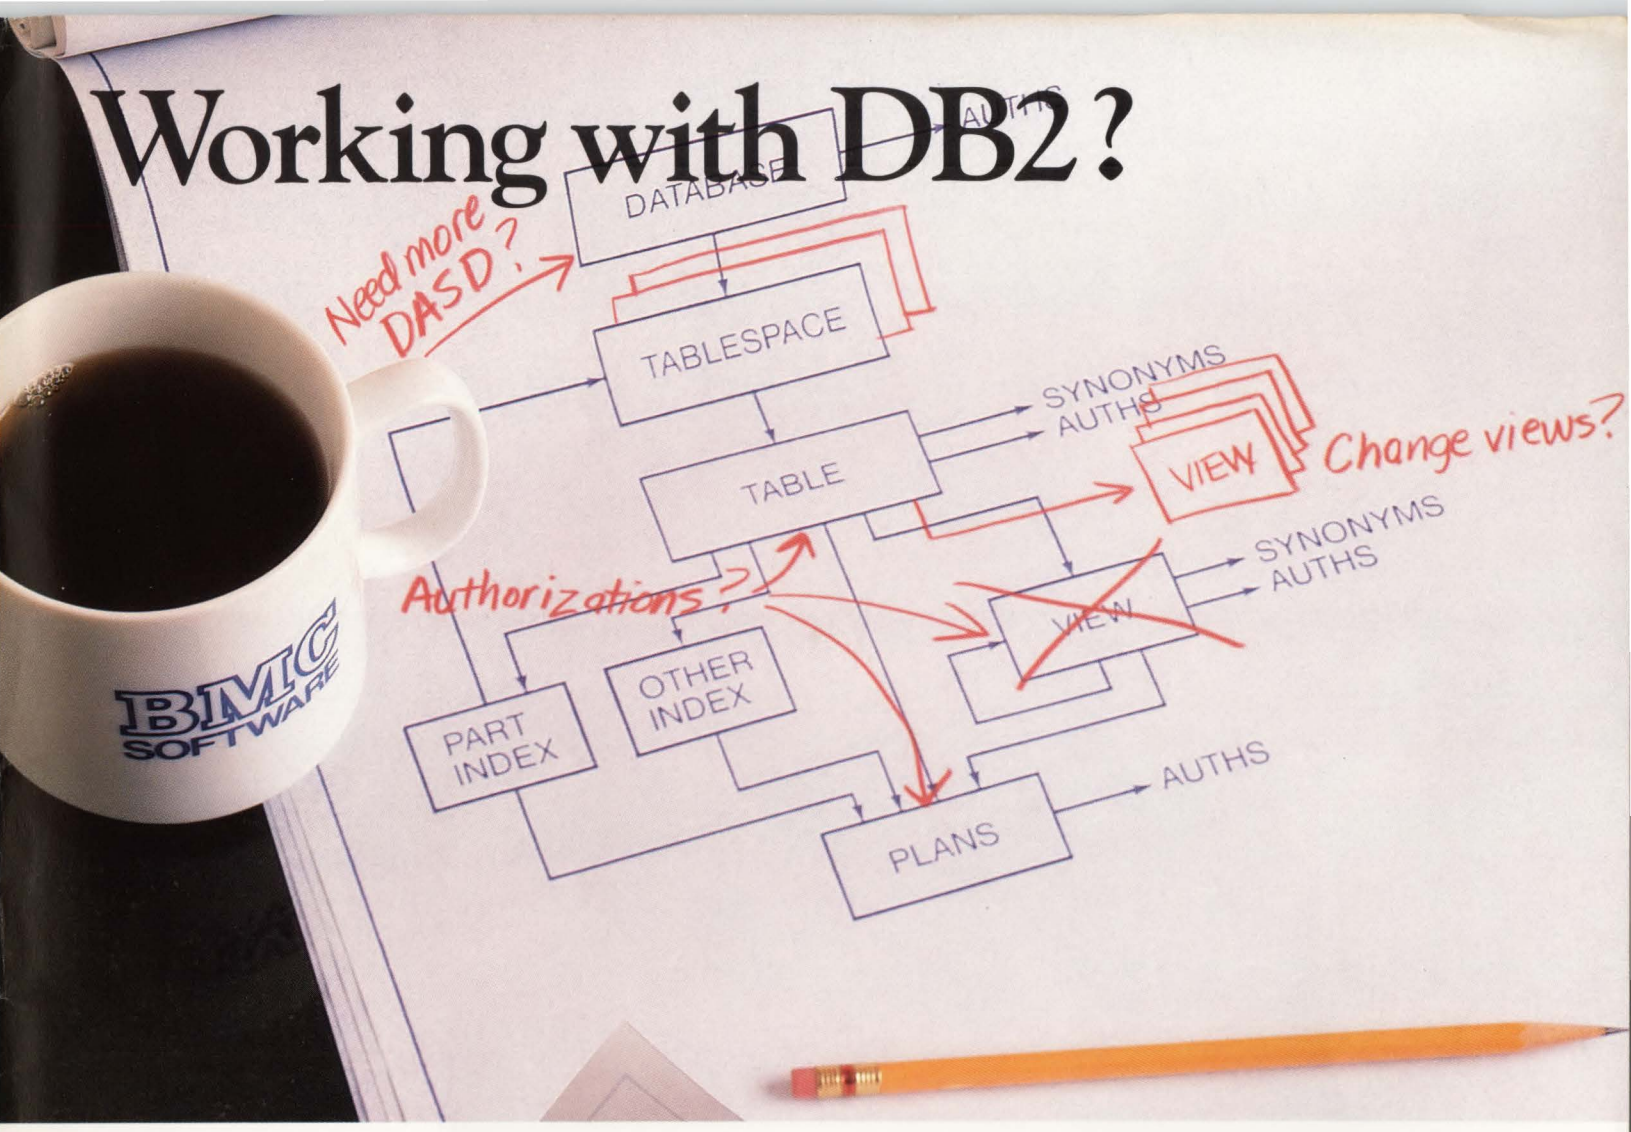

Work with BMC.

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 **DB2 MASTERPLAN**™—a comprehensive set of DB2 products—your work is complemented by a company that has worked extensively with DB2 and knows what you need to keep your system running efficiently.

DB2 ACTIVITY MONITOR—displays and collects real-time data and historical data from MVS, IMS, CICS and DB2; provides more functionality than any other DB2 monitor available.

DB2 MASTERMIND™—complete DB2 administration in one product consisting of:

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 the reader service number. Or, call BMC Software, Inc., **The Complete DB2 Company**™, at **1-800-841-2031**, in Texas, call collect, **713-240-8800**.

BMC SOFTWARE

P.O. Box 2002
Sugar Land, Texas 77487-2002

IBM Business Partner

BMC Software, Inc. also has offices in:
England, France, Italy, Japan and West Germany.

CIRCLE #103 on Reader Service Card ▲

If you are presently using these systems, let IAM® reduce your Nightly Processing Time 50% to 80%

American Management Systems
CUFS (College Financial Systems)

American Software
General Ledger

Cyborg
Payroll

Computer Associates
CA7 – Banking Infopoint – CA11

TCC
Life 70

Data Design Associates
Financial Packages

Cincom
Mantis

HBO & CO
Hospital Applications

Group 1 Systems
Zipcode File

IBM
CICS – SMP/E
RMDS Netview DISOSS

Florida Software
Banking Applications

McCormack & Dodge
General Ledger, Payroll, etc.

Management Science of America
Accounts Receivable, Payable, etc.

Pansophic
Easytrive

Stockholder Systems
PEP+

In House Applications

IAM vs VSAM

IAM REDUCES THE SIZE OF YOUR VSAM FILES BY 30 TO 70%
IAM REDUCES VSAM EXCPs BY 50 TO 80%
IAM REDUCES VSAM CPU TIME BY 20 TO 40%

CICS and Batch applications make heavy use of keyed indexed VSAM (KSDS) files. VSAM is a prime bottleneck to the performance of these systems. IAM, a transparent alternative to VSAM, eliminates the VSAM bottleneck by slashing I/O and CPU time. IAM's Real Time Tuning monitors I/O activity, dynamically changing the number of buffers and I/O chaining. IAM completely eliminates the EXCPs to the VSAM index. IAM executes fewer instructions per record request. Batch processing times are dramatically reduced.

Call for a Free No Obligation
90 Day Trial
Supports MVS, MVS/XA and MVS/ESA
Makers of FDR and ABR

INNOVATION® DATA PROCESSING

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

CIRCLE #97 on Reader Service Card ▲