



MAINFRAME JOURNAL

For Users of IBM System/370 Architecture & Compatible Systems

July/August 1988



In this issue . . .

- Using DB2
- MVS Cross Memory Services
- Multiple Session Management
- Performance Management Primer

- VM Security
- NetView and NetView/PC
- Software Maintenance
- Selecting CASE Technology
and much more. . . .

THE BEST KEPT

SECRETS

ABOUT VM SECURITY...

are yours **FREE** in "What About VM Security?", an informative, comprehensive booklet from VM Software, Inc.

Find out what VM security **really** means. How is it different from security for other operating systems? How can you be sure that **your** system is safe? The answers to these and all your questions about VM security are found in this free booklet from the experts at VM Software.

Topics covered in your **FREE** booklet:

- VM's Provisions For Security
- Password Options
- Data Sharing
- Data Integrity
- Access Rules and How They Work
- The CP Directory
- Recommendations on Making Your System More Secure



800-562-7100

In Virginia call (703) 264-8000

In Canada call (800) 533-5128

WHAT ABOUT VM SECURITY?



I want to know the secrets about VM security. Send me my **FREE** booklet. I understand that I am under no obligation.

Name _____

Title _____

Organization _____

Address _____

City _____

Phone _____

State _____

Zip _____

Fill out this form and return it to VM Software, Inc., 1800 Alexander Bell Dr., Reston, VA 22091. For faster service, call **800-562-7100** in the U.S., **800-533-5128** in Canada, or **703-264-8000**.

1-MFJ-8807

CIRCLE #20 on Reader Service Card



**At level 2,
no one hears
you scream.**

But it doesn't matter anymore.

Now, for the first time, there's Eyewitness™—a *unique* product that helps systems programmers solve the mystery behind CICS system dumps. It's quick, easy, and revolutionary.

At last, you can take system dumps in *less than two minutes*—minimizing the impact on other production regions. Enjoy menu-driven screens. Take advantage of a 50% savings in disk storage space. Rapidly debug dumps with online diagnostics. Instantly analyze storage violations, CICS abends, operator cancels, and snap dumps—**online**.

Eyewitness. There's nothing like it. Best of all, it's another product from the makers of The Monitor for CICS, Landmark Systems Corporation.

For a *free, 30-day trial*, call us today at **1-800-227-8911** or **1-703-922-7101**—and scream no more!

LANDMARK™

eyewitness™
SOLVE THE MYSTERY

Landmark Systems Corporation
6551 Loisdale Court, Springfield, Virginia 22150

CIRCLE #29 on Reader Service Card

CONTENTS

July/August
1988

Volume III
Number 3

DEPARTMENTS

- 6 Publisher's Comments
- 8 Reader Forum
- 67 Advertiser Index
- 86 The Tech Advisor
- 90 Product Update
- 94 Vendor Profile

COVER:

The cover photograph by
Harold Sund from
The Image Bank.



22

Cross Memory Services

How Secure is Your VM System		
	VM SYSTEM	ACCESS SECUR
PRIORITY ISSUE	No	No
Protects data at file level	No	No
Protects data on tape		Yes
Control of resources	Some	No
(control of resources on computers, disks and so on)	No	No
Provides user-level security control	No	No
Provides dependency on passwords	Somewhat	No
Easy to use/implement		No
Assures data is secure in event of system crash	No	Yes
	Some	Yes


40

Securing Data on the
VM System

FEATURES

- 10 Using DB2 — A User's Perspective and Experience *By David Allen Pollock*
- 14 Multiple Session Management — A Cost Effective Way to Increase MIS Productivity *By Lloyd A. Hagemo, Jr.*
- 22 MVS Cross Memory Services — Handle With Care *By Bill Carico*
- 26 SNA Network Management Using NetView and NetView/PC *By Jay Ranade and George C. Sackett*
- 30 Software Maintenance: Undermanaged and Understaffed *By Jan Snyders*
- 35 Systems Application Architecture — An Overview *By Phil Layland*
- 36 Water on the Rock — A Performance Management and Capacity Planning Primer *By R. Douglas Swords*
- 40 Securing Data on the VM System — A Rational Approach *By Steve Halperin and Steven Rev*
- 52 A PROLOG-Based Expert System for Tuning MVS/XA *By Dr. Bernard Domanski*
- 63 Selecting CASE Technology and Software *By Howard W. Miller*
- 68 Apply Large Network Management Strategies to Smaller Networks *By Stephen N. Steinmetz*
- 76 Software Review — Tower Systems' EPIC/VSE and EPIC/CMS Keep DASD Response Times Down *By John Kador*

MAINFRAME JOURNAL © (ISSN 0895-5751) is published bimonthly by Thomas Publications, Inc., 10935 Estate Lane, Suite 375, Dallas, TX 75238, (214) 343-3717. Second class postage paid at Dallas, TX and other locations.
POSTMASTER: Send address changes to: MAINFRAME JOURNAL, P.O. Box 38185, Dallas, TX 75238.



**“There are no holidays
in a hospital.
Or for its CICS.”**

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

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

That's why UCLA chose The Monitor For CICS to manage CICS performance at its computer center. “What initially attracted me to The Monitor was the fact that it made it easy for operators to look at medical application and terminal regions,” remembers **Jerry Johnson, supervising CICS programmer.**

“We're running about 450,000 transactions/day and still maintaining half-second response times.”

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

The Monitor is the *complete* CICS performance management system that'll help you save the day. Become the hero in your CICS community! For a *free, 30-day trial* of The Monitor For CICS, call us today at 1-800-227-8911 or 1-703-922-7101.

LANDMARKTM


**THE
MONITOR[®]**
FOR CICS

MINDS YOUR BUSINESS

Landmark Systems Corporation
6551 Loisdale Court, Springfield, Virginia 22150

CIRCLE #16 on Reader Service Card

Relocating our office was a "moving" experience. While the change of address was well worth the effort, there were a few minor problems like the phones not working and the mail not being delivered. However, all that is a thing of the past and it is now business as usual.



Robert H. Thomas

Multiple Session Management

Lloyd Hagemo believes one of the fastest ways to increase productivity of MIS personnel is by using a multiple session manager. He explains how session management products may be used and how to evaluate the alternatives for providing multiple session support in his article on page 14. At the end of his article is a list of multiple-session manager products giving a brief description of the products and the opportunity to request more information by circling the corresponding product number on the Reader Service Card elsewhere in this issue.

Software Maintenance

Is software maintenance still a problem? Are new tools and methods the answer to solving the problem? Is there, in fact, a solution to the maintenance problem? Jan Snyders asks nine industry leaders for their opinions on these questions and others in her article on page 30.

SAA

Phil Layland in his article on Systems Application Architecture (SAA) gives an overview of what he feels may be the most significant IBM announcement in a decade. While SAA is still just a concept, understanding it now will enhance future decision making. Layland describes SAA, the reason for its existence and recent inclusions in his article on page 35.

Is Your Data Secure?

In "Securing Data on the VM System" by Steve Halperin and Steven Rev, the authors address the fact that VM allows opportunities for sensitive data to be exposed, altered and/or destroyed. Their solution? See their article on page 40.


PROLOG-Based Expert System

Dr. Bernard Domanski sheds new light on the application of Artificial Intelligence (AI) to Computer Performance Evaluation (CPE). In his article on page 52, he presents some of the issues involved in building a PROLOG-based expert system for tuning IBM's MVS/XA operating system.

Network Management Strategies

Effective network management is as vital to small companies as it is to larger ones says Stephen Steinmetz in his article on network management strategies. He maintains that by adapting strategies used to manage large networks, a company with a smaller network can provide better service to users of its current network. See if his strategies starting on page 68 can work for you.

Network Support Organization Evolution



Publisher and Editor-in-Chief

Robert H. Thomas

Associate Publisher

Martha Thomas

Editor

Carol M. Hoag

Copy Editors

Judy Beller

Pat Warner

Art Director

Laura Page

Illustrator

Jeff Hill

Advertising Manager

Denise Thomas

Marketing Services

Sally Webb

Circulation Manager

Janice Porter

Assistant Circulation Manager

Nancy Crawford

Administrative Manager

Marian Davenport

Typography Coordinator

Dub Pope

Jaggars Chiles Stovall

Printing Coordinator

Joe Jarzen

PennWell Printing Company

BPA Membership

Applied For April, 1988.

SUBSCRIPTION RATES & INQUIRIES: Subscriptions are free within the United States. One-year subscriptions in Canada and Mexico are \$30; all other foreign, \$48. All subscriptions, remittances, requests and changes of address should be sent to MAINFRAME JOURNAL at P.O. Box 38185, Dallas, TX 75238, (214) 343-3717.

MAINFRAME JOURNAL is published independent of IBM. Opinions expressed in MAINFRAME JOURNAL are not to be considered official expressions of IBM.

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


The Secret To Maximizing Your 4381's Performance Lies In The Cards.

Main storage upgrades are the most reliable and cost-effective way to improve the performance of your 4381 computer. But only EMC main storage upgrades let you *maximize* the power of your 4381 computer system.

That's because only EMC provides main storage upgrades which perform memory writes twice as fast as IBM's main storage cards.

EMC's Main Storage Cards Run Cycles Around IBM's Cards. How do we do it? Simple. Our upgrades use newer technology to perform a memory write in 11 processor cycles. It takes IBM's upgrade 22 cycles to perform the same memory write. This means our upgrades increase the availability of main storage to your 4381's processor and take advantage of otherwise wasted CPU cycles.

With EMC main storage upgrades you avoid the expense of a model upgrade while still getting the performance you need to meet increasing demands on your computing resources.



Maximized Performance is Only the Beginning.

As the world's largest independent manufacturer of main storage upgrades for IBM's mid-range computers, we deliver a lot more than just improved performance.

Start with reliability. All our upgrades for the 4381 undergo a 100-hour test and burn-in procedure, including qualification in a 4381 computer.

Then there's your \$20,000 savings on every 16MB upgrade, and a guarantee that use of our upgrades will have no effect on your IBM maintenance service.

Plus our coterminous leases, immediate delivery and responsive service programs.

To find out more about EMC's high-performance main storage upgrades for your 4381 computer, call today or write: EMC Corporation, Hopkinton, MA 01748-9130.

For more information, call today:

1-800-222-EMC2

(In Mass., 617-435-2541)

European Headquarters call: 01-6685511

In Toronto call: 416-368-4726

In Vancouver call: 604-662-3911

IBM is a registered trademark of International Business Machines Corp.
Copyright EMC Corp. 1986

EMC²

The System Enhancement Company.

CIRCLE #95 on Reader Service Card



HELP

We have a question on file integrity that we are hoping *MAINFRAME JOURNAL* readers might have an answer to.

We run a DOS/VSE shop under VM in which we have three DOS machines running. Most of our VSAM files are specified as share option 4 which allows us to update from either batch or CICS simultaneously at any time of day as needed. The problem, of course, is that you can't do that across machines without losing file integrity. By specifying SHR on a disk drive at IPL time, VSAM will do I/O to the DOS lock file and protect the file, but that stops any access from another machine until the file is closed in the first machine. Does anyone have a good way to do this without closing files or going to MVS?

*Kenneth O. Marquardt
Paper Converting Machine Co.*

This letter is a request for help from readers of *MAINFRAME JOURNAL*. Our production machine is a VSE/SP 2.1.5 guest.

My problem is this. Because of numerous personnel changes and the natural antipathy of DP personnel to do the dull house-keeping chores, we have a private library full of obsolete phases which no one knows anything about. I move this stuff into a temporary library out of the search chain. If a job blows up, the phase is easily moved back to its old home. Eventually, I backup the whole sublibrary to tape and delete it. This system works, but that's all I'll say for it. What I would prefer is a system that tags phases with the date last executed.

Does anyone out there have or know of software that does this? I would greatly appreciate hearing from your readers who have found a solution to the problem of obsolete mystery phases.

*Warren Jackson
Washoe County School District*

Editor's Note: Please send any responses and pleas for help to: Reader Forum, MAINFRAME JOURNAL, PO Box 38185, Dallas, TX 75238.

TECH ADVISOR RESPONSES

I would like to amplify your answer to the first question in the January/February 1988 issue concerning CICS HANDLE CONDITIONS. We code a single HANDLE CONDITION at the beginning of a CICS program for processing efficiency.

To allow for the situation where control is transferred to a single routine as a result of a file condition (NOTOPEN, NOTFND, etc.) which could arise from any of several different files, we simply check the EIBDS field of the Exec Interface Block at the beginning of the generic routine pointed to by the HANDLE CONDITION. This 8-byte field contains the name of the last dataset accessed and can be incorporated into an error message to be sent back to the terminal, compared to a literal in an IF statement, or whatever.

We sometimes test other EIB fields along with EIBDS, depending on our processing requirements. This solution is easy to implement and does not require an upgrade to CICS 1.7.

*Brian D. Phillips, CDP
General Dynamics Corp.*

In your March/April 1988 issue the question was raised as how to move a production file when the JCL is maintained in SLI books. The problem seemed to be how to make mass changes when there is too much JCL to punch onto a single tape. You suggested one possibility, nesting SLIs which contain only DLBL and EXTENT cards. Since some JCL, especially CICS start-up JCL, tends to live in the POWER reader queue, this approach is not good. SLI cards are resolved as the cards are read by POWER, not at execute time. Thus, even though you changed a SLI book for the new file placement, any JCL previously submitted to POWER would be unaffected.

There is another solution. Place your DLBL and EXTENT cards in PROCs. PROCs are resolved at execute time, not job submission time. Thus, you would code '// EXEC PROC =PAYMAST' in your JCL where you would have coded the DLBL and EXTENT cards for file PAYMAST in the past.

Now, simply change the PROC when you want to alter file placement.

VSE/SP provides additional capabilities. Operands could allow you to use the same PROC and yet specify override parameters such as file name and VSAM disposition. For example, a program which loads file PAYMAST might need to reference the file by filename PAYLOAD with an open disposition of NEW on the VSAM DLBL card rather than the normal OLD. Instead of coding the DLBL and EXTENT cards special, just code '// EXEC PROC =PAYMAST,NAME =PAYLOAD,OPEN =NEW'.

With these techniques, the DBA can move files as often as suits his need without concern for JCL. JCL need never be modified simply because a file has been moved. Processing of the PROCs might require a few additional disk I/Os, but that is a small price to pay for the increased flexibility gained.

*James E. Holler
Little Blue Systems, Inc.*

I was just going through my back issues of *MAINFRAME JOURNAL* looking for information I remember reading on how to get CICS programs loaded end-to-end rather than page aligned, when I read your reply to a user's question on how to maintain JCL contained in SLI books (March/April 1988 issue, page 80 — last question). Your suggestions are great. We do not keep DLBL/EXTENT information in SLIs or procedures but rather in standard labels or partition standard labels. This way I have only one place to make changes (something I haven't had to do in a long time). By using partition standard labels, I can also limit a particular application to a specific partition. It sounds almost as if this user is constantly moving files, and if that is the case I would encourage them to sit down and do some serious DASD planning.


One additional method the user may want to try is to execute his COBOL conversion program letting SLI supply the input rather than punch it all to tape first; this method may be a little faster. This of course will not provide the CATALS/BKEND statements as SSERV does, but that could be compensated for by modifying the COBOL program to look for a book header statement and use that statement to create the appropriate CATALS/BKEND statements. The resultant punch output should be spooled as DISP=I to execute as a catalog job following the conversion run. In order to place some JCL around the converted books, our friend could use DITTO, OBJMAINT or some other utility.

There are a couple of things to be aware of when using this method. The first thing is the possibility of filling the POWER queue file (as you mentioned) because of the volume of data involved. Secondly, if you choose to write output to VSAM managed SAM instead, you cannot use OBJMAINT (I don't know why, IBM just says you can't). I prefer spooling output back into the reader. If the volume of data threatens filling the queue, you can break the process down into several smaller runs.

I have found in some cases, CICS command translation in particular, that a blank card image is spooled at the end of job (when the punch is closed) which shows up as an erroneous statement during assembly time. This is due to a peculiarity in POWER when the punch device is defined as a 3525 and spooled with DISP=I. This extra blank "card" can be eliminated by spooling DISP=I punch output through a punch which is defined as a 1442-N2.

Well, here is my two cents worth on the subject; now back to my search. Keep up the excellent work. We all benefit from your efforts. I, personally, have learned very much by reading *MAINFRAME JOURNAL* and once in a while now, I find that I have some ideas to offer. Thanks again.

*Jim Wilson
San Mateo Union High School District
See Forum page 80*



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

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

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

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


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

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

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

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

Westinghouse
SOFTWARE SOLUTIONS



CIRCLE #128 on Reader Service Card

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

Using DB2

A user's perspective and experience

By David Allen Pollock

A basic understanding of relational theory is an extremely necessary prerequisite to the proper use or evaluation of a relational product such as DB2. Fortunately, there is a certain intuitive quality to the underpinnings of the theory making it easier for both data processing and non-data processing professionals to understand.

Relational theory was first proposed as the basis for a DBMS by Dr. E. F. Codd in 1969 and was subsequently extended in 1979 (it should be noted that the extensions represent applications of the original theory rather than revisions to it). A full description of relational theory is quite involved and beyond our scope. However, the basics necessary to appreciate the power and simplicity of relational DBMSs is easily presented.

In relational systems, data is perceived as a collection of tables. It is important to note that this is not an indication of how the data must be stored, but rather how the user perceives that data is stored, retrieved and updated. Data is not processed a record at a time but is accessed in sets. In other words, many rows and columns of data may be addressed in one command. This is a powerful feature of the relational model which can be both under-estimated and under-used by record-at-a-time programmers. The operators for retrieving and updating data in the tables also have the quality of closure: the result of any operation on a table(s) is always another table. This is an important feature in that it allows for the nesting of operations, giving the model the power to handle virtually any logical expression within the limits of first order predicate calculus.

Before discussing the operators, it is important to review the necessary com-

***DB2 is a powerful,
relational DBMS.
Although that statement
is in itself redundant,
it is important to note
that there are some
facets to the use of a
relational DBMS other
than just the fact
that it is relational.***

David Allen Pollock

ponents of a table. Every table has a title for the table, columns which also have headings and rows of values in each column. For example, a table might be called COURSES having the columns Course#, Course Name and Dept# (see Figure 1). The rows are distinguished and identified by a primary key: one or more columns whose values are unique within the table (in this case, Course# in the COURSES table and Dept# in the DEPARTMENTS table). While other columns may also uniquely identify the rows in a table, there is only one primary key. Other important aspects of table definition (such as domains) should be studied in detail by table designers.

There are several different operators in relational theory, but the most important to be aware of are restrict, project and join. Restriction is when a subset of rows in the original table is desired. For example, when information from the COURSE table is wanted for only department 111, the query result can be re-

stricted to include only those rows which have the value 111 in the Dept# column (see Figure 2). Restrictions may be made on an equivalence test, a "likeness" test or any other Boolean type of comparison.

Projection is the process of selecting which of the available columns are to be included in the query result table. For example, if only course numbers and names are required (as in student transcript preparation), then the query result could include only the desired columns: course number and course name (see Figure 3).

Finally, join is the most powerful and the most dangerous of the operators. It allows for the combination of data in multiple tables into a single result table. Remember that, since all data is perceived only as a collection of tables, then it is necessary that the combination of two or more tables be stated only by referencing data within those tables (rather than by following a predefined and separately maintained set of pointers). Generally, one of two tables to be joined together will contain a column which is the "primary key" of another table. Such a column is referred to as a "foreign key." The join is performed by matching all of the rows in the two tables where the values in the foreign key column in the first table are equal to the values in the primary key column of the second table. For example, if I wished to obtain the course number, department number and department name, I must access two tables where the department number in the course table (a foreign key) is equal to the department number in the department table (a primary key — see Figure 4).

Other operations for creating, updating and deleting rows from tables are similar to restriction in that they also operate at the set level on qualifying rows. Despite the fact that so much attention is paid to



Before your 308X goes down, call Sorbus.SM

Keep little problems from becoming big ones, with Sorbus service for your IBM® 308X.

It includes all of our unusual Sorbus usuals: Our two-hour response time, our 230,000 part-number inventory (more than 6.2 million parts, last time we looked), and our convenient scheduled preventive maintenance.

Not to mention our reputation—among people like the readers of *Data Communications*, who've voted us "Best Service Organization" and, for nine years running, the one having the "Best Price/Performance Ratio." And like the readers of *Datamation* and *Computer Decisions*, who've rated us the Number One independent computer maintenance company consecutively for eleven and eight years, respectively.

We already take care of computers and peripherals at 60,000 sites—including more IBM equipment than anybody but IBM.

We'd like to take care of your 308X. Before it's too late.
Call us at **1-800-FOR-INFO.**

SorbusSM
A Bell AtlanticTM Company

50 E. Swedesford Road
Frazer, PA 19355

Sorbus is a service mark of Sorbus Inc.

IBM is a registered trademark of International Business Machines Corporation.

CIRCLE #93 on Reader Service Card

the flexibility of retrieving information, one should not infer that relational systems are only for query. Their power in updating potentially thousands of rows with a single command is staggering.

FIGURE 1

COURSES		
Course#	Course Name	Dept#
102	Biology	111
247	Calculus	222
394	Algebra	111
410	Greece	333

DEPARTMENTS	
Dept#	Dept Name
111	Science
222	Math
333	History

Sample Tables

FIGURE 2

The operation:
 Restrict COURSES
 where Dept# = 111
 yields:

COURSES		
Course#	Course Name	Dept#
102	Biology	111
394	Algebra	111

Restriction

Two important points of note are brought out here. The first is that the operators may be combined in any way which the user desires (Figure 4 actually uses a projection in addition to a join). The other is that the values of all foreign keys must correspond to those of primary keys. This is the problem commonly referred to as referential integrity (which has always existed, but which has gained more publicity with relational systems).

Having presented the data as tables to the user, who may remain unconcerned about the underlying physical data structures, a level of physical data independence is achieved. This means that the DBA may add or remove indexes and/or restructure the data without impacting the work of the users. This is a crucial productivity and performance aspect of the relational model.

These are the basic, yet important points to know about relational theory. They require a reorientation to thinking in terms of data, not pointers, and sets instead of records. Often users are better able to deal with this change in mindset than are applications developers, because it is intuitively how they (the users) have been perceiving their information all along.

DB2: A Relational DBMS

Having addressed some of the important aspects of relational theory, it is important to consider how a DBMS becomes relational. Generally, the degree of support of the theory would determine the degree to which the DBMS could be considered relational. It should be noted that no DBMS currently supports all of the relational model or theory which Codd has defined. However, some do satisfy it more fully than do others. DB2 does quite well in some areas and has a ways to go in others.

Within DB2, all data is indeed created, stored and manipulated as though it is in a tabular format. Through one comprehensive language, SQL, most of the operations are supported in one way or another. Restrict, project and join are provided in a single command structure called the SQL Select command (in essence, the SQL Select-From-Where structure corresponds to the relational Project-Join-Restrict operators respectively). However, DB2 has no knowledge or capacity for identifying primary and foreign keys. This unfortunately leaves the enforcement of referential integrity to the users, a wearisome and error-prone task.

A more basic point is that, since DB2 does not even contain the information as to which columns constitute the key fields, that knowledge must reside outside of the system, despite the reliance of almost every operation on such knowledge. When joins are performed, it is necessary that the DBMS know which field in each table to compare (primary and foreign keys). DB2, lacking such knowledge, will in fact build a cartesian product of all rows in both tables, counting on the user to also provide a restriction indicating which rows of the result table should be retained (in reality, this restriction is often performed prior to the building of the result table for the sake of efficiency). Uniqueness of columns is enforced through the creation of unique indexes, whereas this should be handled as a part of the table creation and key indication. In addition, DB2 does not provide the ability for indicating underlying domains on which columns are based. This also must be recorded outside of the DBMS proper.

DB2 also supports a limited amount of physical data independence in a variety of ways. Indexes on tables may be created and dropped by the DBA as necessary to improve performance. These indexes are or are not used in the solution of queries as determined by the query optimizer.

Also, columns may be added to tables without impacting existing queries (or requiring larger I/O areas in existing programs). And, through the use of views (which are essentially stored queries), some additional independence is obtained. However, changes to the underlying table structure still have a direct effect on queries which perform updates and should be done carefully.

In spite of these limitations, however, DB2 is a robust, full feature DBMS which provides ample power and flexibility for real production applications. With the emergence of a SQL standard, the language for manipulating data within relational DBMSs will become more similar across products. This implies that the success or failure of the DBMS in the market will hinge even more than ever on its ability to perform.

FIGURE 3

The operation:
 Project from
 yields:

COURSES	
Course#	Course Name
102	Biology
247	Calculus
394	Algebra
410	History

Projection

FIGURE 4

The operation:
 Join Courses, Departments
 yields:

DEPARTMENT-COURSES		
Course#	DEPT#	DEPT NAME
102	111	Science
247	222	Math
394	111	Science
410	333	History

Join

DB2: A User's Experience

We have designed DB2 tables with a variety of sizes and shapes. Some tables have many columns, some only a few. Some tables have only a handful of rows while others are in the tens of millions. The most interesting aspect has been, however, that we have dismissed entirely the notion that DB2 is for read only, or that it is not for production. Rather, most of our systems include on-line updates in production. The fact that most of our

See DB2 page 85



To solve terminal user problems fast, you need more than the next best thing to being there.

If your biggest hang up is time lost while solving users' problems by phone, you need MULTITERM/VM™, the VM session manager, or MULTITERM/SNA™ the ACF/VTAM session manager for VSE systems. Now, problems that users experience with systems or applications can be brought straight to you. That's because MULTITERM/VM & MULTITERM/SNA can display them on your own terminal screen. No more frustrating phone calls. No more running around.

You can work on a user's session without leaving your desk, whether the problem's in the next room, the next state or the next country. That's because MULTITERM/VM and MULTITERM/SNA feature Help Desk facilities. So you can take a session that a user is having problems with and display it on your own screen. Then, after you've solved the problem, you can give the session back to the user. Or you can monitor the screens being displayed on another user's terminal at your terminal. The Help Desk facilities also enable you and another user to jointly enter commands on the same screen display from anywhere.

That's not all. MULTITERM/VM and MULTITERM/SNA feature instant session switching and message broadcasting capabilities. And, operating in conjunction with each other, the products allow CMS users to access VTAM applications and vice versa.

You'll also be glad to know that all our technical support is located in the U.S. and that we provide site licensing as well as purchasing and leasing options.

All these features, together with comprehensive user profile facilities make MULTITERM/VM and MULTITERM/SNA the ideal session manager for VM and VSE systems.

So contact BlueLine today and find out how MULTITERM/VM and MULTITERM/SNA will put you face to face with problem resolution. Just **call 800-826-0313**. In Canada and Minnesota **call 612-542-1072**. *Now that's a good reason to use the phone!*

MULTITERM/VM & MULTITERM/SNA

BlueLine
SOFTWARE INC.

New Advances in Performance, Productivity and Planning.

CIRCLE #48 on Reader Service Card

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

Multiple Session Management

A cost effective way to increase MIS productivity

By Lloyd A. Hagemo, Jr.

Everywhere you turn today, you hear about a company that has done something innovative to increase the productivity of its employees. This desire to boost productivity is felt in all departments of a business from the assembly line to the executive suite. This includes the MIS department. One of the quickest and easiest ways to increase the productivity of the MIS personnel is to give them a multiple session manager.

By enabling users to run multiple sessions concurrently, a session manager eliminates needless and time consuming logoffs and logons when switching from one system to another. Additional features of session manager products enable operations and systems personnel to respond to users' problems with a minimum of disruption to their own work. This article explains the uses of session management products and describes some of the things that you might consider when evaluating the alternatives for providing multiple session support.

Networks and distributed processing have resulted in complex environments for users. With several on-line systems available, it is common for users to access several different systems in the course of a day. Quick access to the information stored in these diverse systems is critical to the smooth operation and success of most businesses. As a result, companies cannot afford the additional time and system overhead associated with logging on and

off when switching from one application to another. The requirement to access several applications concurrently from one physical terminal has led to the development of hardware and software multiple session manager products.

To illustrate how a session manager functions, look at a hypothetical scenario. Assume you are doing applications programming on an on-line CICS development project. The CICS command level COBOL programs are written on TSO, compiled in batch and tested on CICS using Execution Diagnostic Facility (EDF). To cut down on the need for hard copy output, you have Spool Display Search Facility (SDSF) available on TSO to review listings on-line.

The first thing you do in the morning is log on to TSO and complete some changes to your COBOL program. When you are ready, you submit the source for a compile and review the output using SDSF. At this time you have already done some testing, but a few problems remain. If you have a clean compile, you log off of TSO and log on to CICS to continue the testing process.

After NEWCOPYing the program and starting CICS EDF, you key in the CICS transaction. But you find that some of the fields on the map are not being edited properly. After committing the problems to memory, you log off of CICS and log on to TSO to make the necessary changes to your code, hoping you have remembered the nature of the problems. You then

recompile the program and check the output using SDSF. When you are ready to test again, you log off of TSO and return to CICS.

You continue switching back and forth between TSO and CICS all day, rewriting and testing your code. If you get to a point where you need a listing to continue testing, you work on another task while you wait for the listing. When the source listing is available, you return to the job of testing the program.

During the course of the project, you spend a lot of time logging off and on and waiting for listings. That time could be better spent working on the actual code. These needless delays are costing the company extra development dollars and they are holding up the delivery of the on-line system.

Adding a Session Manager

Now add a session manager to this scenario. The session manager provides concurrent access to TSO and CICS. In the morning you log on to the session manager and select TSO from the menu that was tailored to your specific needs. You then make the changes to the COBOL program, compile it in batch and check the listing using SDSF. So far, there is not much difference from your routine before the session manager was installed.

If the compile is clean, you return to the session manager menu by using the menu return key and you select CICS. Figure 1 contains an example of how a session manager's menu might look.

After logging on to CICS, you NEWCOPY the program, start CICS EDF and key in the transaction to start testing. If you find an error, instead of logging off of CICS, you toggle to the TSO session and make the changes, compile your programs and check the output using SDSF. A toggle is a single keystroke that will switch you from one session to another. Toggling back to CICS, you NEWCOPY the program and start testing again. If you forget the exact problem while making the changes to the program, you can return to CICS with a single keystroke (toggle) and the screen with the problem reappears on your terminal.

The ability to return to the test and see the incorrect results while making changes saves a lot of time. As you proceed with your testing, the task of debugging goes much more quickly. You can view the listings on SDSF on one session while testing in CICS using EDF on the other session. Having the listing available a keystroke away on one session while testing on another session eliminates the need for a hard copy listing of the program. You seldom have to wait for a listing and switch your train of thought to another task. Your focus and concentration are on one work item. The result is that you can develop a better system in less time.

Using this example, take a quick look at the kind of savings a company can easily realize. Assume it takes one minute to log off of CICS and log on to TSO. When you log on to TSO, you are either in an edit session looking at your source or SDSF looking at your listing. On an average day, you might switch between CICS and TSO ten times. This means that you lose ten minutes each day on nonproductive logons and logoffs. This does not include the lost productivity that results from not being able to review your test results as you work on a problem or from waiting for hard copy output.

This comes to 50 minutes a week, or 2,500 minutes a year, figuring 50 working weeks in a year. For one employee, this amounts to 42 hours a year. If you multiply this over a staff of ten, you accumulate ten-and-a-half weeks of measurable lost productivity. If you apply this savings over a user base of 400 to 600, the benefits are even greater. Even at the conservative estimate with ten employees, it does not take long to justify a session manager product.

Session Management Alternatives

There are two kinds of session manager products on the market today. IBM and

most 3270 OEM vendors offer hardware solutions and there are numerous software products available on the market.

Hardware Solution

You could solve the productivity problem with terminals and controllers that are capable of running multiple sessions. Because these terminals function independently of host software, this is a good solution when reliability is critical.

Operations and systems personnel are good candidates for this solution. But for the general user, there are drawbacks.

Because your company might have a large investment in 3270 terminals, you might not want to replace your current 3270 base or add multiple session terminals to it. The multiple session terminals cost more than standard terminals and some of these terminals require several

FIGURE 1

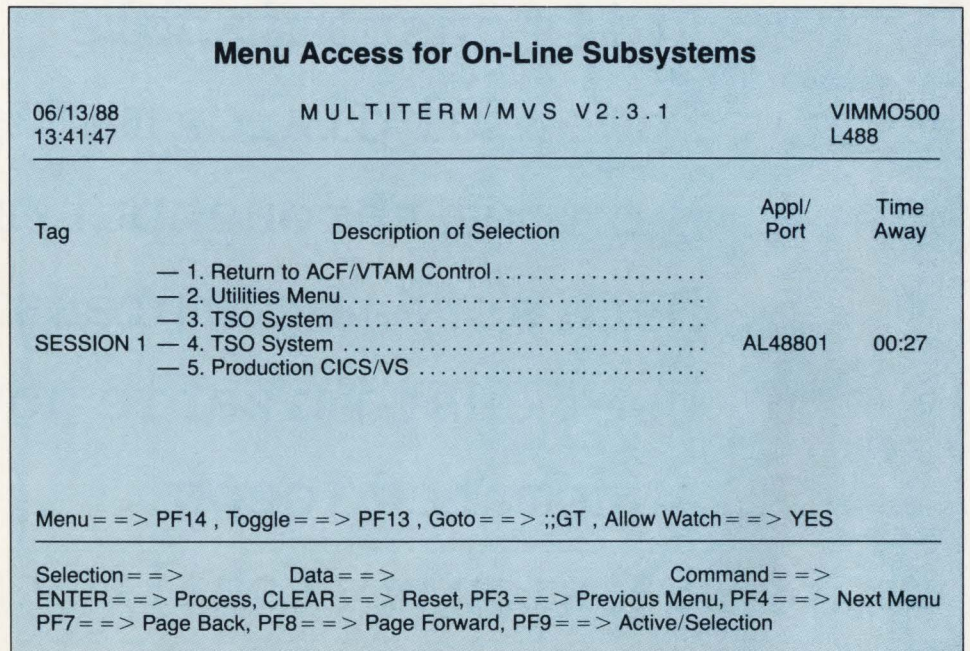
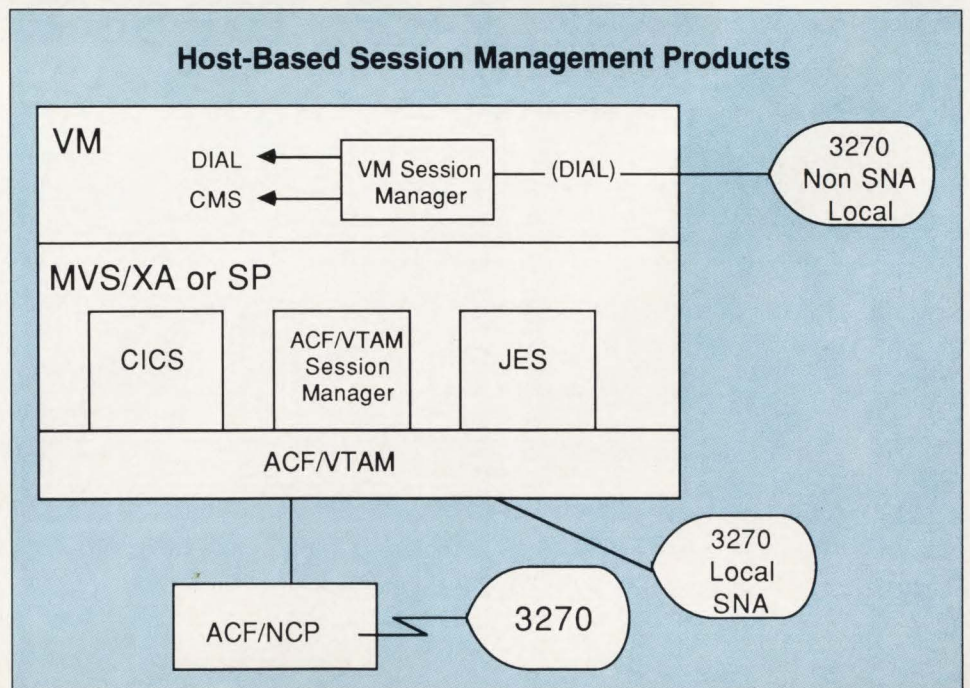


FIGURE 2



Everything You Ever Wanted in a Multi-Session Manager.

1. Unlimited sessions and users
2. Fully customizable
3. Powerful dialog manager
4. Maximum throughput/minimum CPU
5. Data stream compression
6. User customizable hot keys
7. Automated logon
8. Session portability
9. Help desk session view facility
10. ACF2, RACF, TOP SECRET interface
11. Network-wide broadcasting

Network concerns are always the same—optimize resources and maximize flexibility to meet the needs of an ever-expanding user base. CL/SUPERSESSION™ from Candle Corporation is designed to do precisely that.

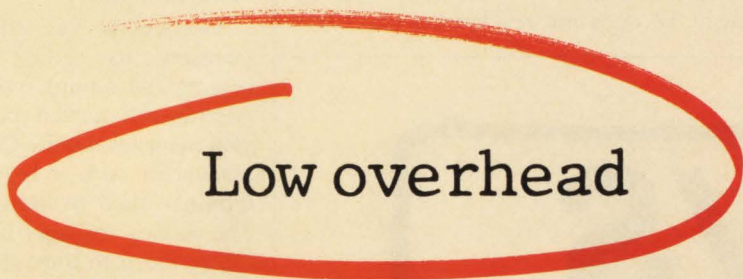
The secret—a powerful dialog manager and procedural language. Many data centers find CL/SUPERSESSION's "off the shelf" facilities a perfect fit for their network. Just as it is. But if you need more, the dialog manager is so easy to use you can customize screen formats and applications, as well as processing logic.

Users can opt for more tutorial assistance. While systems programmers can define their own profiles and access privileges.

CL/SUPERSESSION can be reconfigured dynamically, so you can add users, update network profiles, and change dialog definitions, panels, and exits whenever you need to. Without disrupting service.

And performance is never sacrificed. CL/SUPERSESSION takes full advantage of multiprocessors and MVS/XA support.

And Less.



Low overhead

Most control block and buffer storage can be allocated in the extended private area. And, by using features like VTAM Fast Path and a proprietary storage management technique, CL/SUPERSESSION is the most efficient multi-session manager available.

For your own information folder with all the details on CL/SUPERSESSION, call Terry Forbes today at (312) 954-3888.

When you can have everything ...
why settle for anything less?

Candle[®]
innovators in illumination

1999 Bundy Drive
Los Angeles, CA 90025

Terry Forbes
(312) 954-3888

CIRCLE #116 on Reader Service Card

RACF is a trademark of IBM.
ACF2 and TOP SECRET are trademarks of Computer Associates International, Inc.
Copyright © 1988 Candle Corporation. All rights reserved.

ports on the controller to provide multiple session support. You might find that this hardware costs more per user to implement than a software solution.

Another drawback is that the hardware solution does not provide any additional features besides multiple session support. Software products today offer long lists of features as we will see when we talk about the software solution. I feel a hardware multiple session product has a place in every company, but software session managers offer more features and flexibility.

Software Solutions

Two types of host-based software session managers are found in the software marketplace today. The most common type of software session manager is the ACF/VTAM application that uses the ACF/VTAM API interface to enable a user to switch between SNA subsystems like CICS and TSO. The second type is an application that runs in its own VM service machine. It uses VM logical-device support to establish several concurrent sessions with CMS or to dial into other

service machines like VSE or MVS to gain access to subsystems such as CICS. Figure 2 illustrates the location of these two types of products. Both of these solutions offer several features and they both enhance a company's current terminal base. But before we consider the benefits, take an objective look at some of the negative side effects.

Host-based session managers cause three side effects. These products require memory, they consume CPU cycles and they increase the user's mainframe path length.

The memory is used to manage the buffer images for the user sessions that are currently active through the session manager. Two techniques are employed for terminal image management. The read buffer command can be used to retrieve the image of the terminal from the terminal controller. This solution has several limitations. Graphic and color attributes are not returned to the host on the read buffer command. The result is that when the user returns to a session after toggling away, he does not receive the exact ter-

terminal image he had before he toggled. In addition to the loss of attribute characters, the read buffer technology places an extra load on your network. Transferring large data streams through the network each time a user switches from one session to another is an unacceptable deficiency in today's competitive session manager market.

The best approach to 3270 screen image management is to use the second technique. This simulates the 3274 terminal controller in the host and maintains the buffer image, including all the attributes, in memory. There are no attribute characteristics lost when the user switches from one session to another, nor is an extra burden placed on the network by the buffer images being transferred from the controller. This technique also eliminates the delays associated with waiting for the read buffer command to complete prior to switching the user to the other session. But in both techniques, terminal images are maintained for all sessions and host memory is used.

The second side effect is the session manager's use of CPU cycles when processing and routing input to applications. All transactions that pass through the session manager require CPU time and a fair amount of code is required to direct the terminal input to the proper destination. The session manager developer must optimize the path through the session manager to reduce the overhead and minimize the CPU utilization. The best techniques to accomplish this are to optimize assembler routines used to manage the flow of data through the session manager and to eliminate IBM system service calls whenever possible.

The third side effect is increased path length. Normally a user logs directly on to the destination application. For example, a user logs on to TSO. The host path includes ACF/VTAM and the user's TSO region. When a session manager is involved, the user logs on to the session manager and the session manager logs on to TSO. The path includes ACF/VTAM, the session manager, ACF/VTAM (again) and the user's TSO region. It is critical that the session manager provide the shortest, most direct code path possible when passing inputs to outputs.

You should consider the above side effects when you evaluate a software session manager. But do not let them deter you from also considering the benefits of a software solution for session management. The many features and benefits far

MVS

PERFORMANCE ENGINEERING

- * We Guarantee performance improvements!
- * If our engineers cannot improve your systems performance according to *exact written objectives* - you pay nothing!
- * Done without any additional hardware or software.
- * Improve response times, maximize batch throughput, and defer hardware purchases.
- * Save potentially millions on this years budget.
- * Initial system analysis is absolutely free of charge or obligation.

**IF WE DON'T SAVE YOU MONEY
WE CHARGE NOTHING!**

Phone: 1-800-562-8722

(In Texas: (713) 451-0191)

Bennett Software Inc.

Engineering Services

POB 96694, Houston, TX 77213-6694

CIRCLE #139 on Reader Service Card

outweigh the disadvantages of these side effects.

Features and Benefits of Software Session Managers

Turning to the positive side of software session managers, I will describe the benefits that session managers provide in addition to multiple session support. The productivity gained by being able to run multiple sessions is a given benefit of all session manager products. The other categories of benefits that I will discuss are access security, help desk services and network performance.

Access Security Features

Session manager products are not security packages. System security is a complex topic and most installations' requirements go far beyond the features provided by session managers. Nevertheless, there are many log-on processing features that relate to security. I have included all of them in this access security category.

User Profiles

Profiles are used by session manager products to identify the characteristics of users. Profiles can define which menu the user will see, how many sessions the user can have and the virtual terminals that are available for this user. In addition, the user's control keys for such things as toggling, issuing commands and returning to the menu are defined in the profile.

Passwords and Exit Points

When a user logs on to a session manager, he is usually presented with a security screen that requests him to enter a user code and password. The user must enter these before he is given access to the session manager. Exit points are provided in most session managers to enable the systems staff to interface to their security products and verify the user's access. In addition, most products provide a way to establish global, system level profiles that make the definition process easier for installations that are not as concerned with security.

Session Follow and Reconnect

Session follow and reconnect are two other features that are associated with log-on processing. The session follow feature allows a user to leave his office, go to another terminal and have all the sessions associated with his user identification moved from the office terminal to the new location when he logs on.

This capability would be very helpful to a systems programmer. Many times the systems programmer must go to see a user to help him with a problem. The systems programmer leaves his TSO identification logged on in his office. He arrives at the user's terminal, starts working on the problem and finds that the user's TSO identification does not include the access authorization that the systems programmer needs to solve the problem. The systems programmer must return to his office and log off his TSO identification before he can continue to work on the problem with the user.

If he had a session manager with the session follow feature, the systems programmer could log on to the session manager from the user's terminal with his own session manager profile. All the sessions that he had in his office including his TSO identification would transfer to the user's terminal. This would allow the systems programmer to work on the problem without returning to his office at all.

Reconnect is like session follow, with the exception that the user is not currently active when he logs on. The user has logged off or has disconnected from the session manager and the session manager has kept all of his sessions active until he

returns. Upon the user's return, reconnect allows him to pick up the existing sessions and continue working without logging on to several subsystems again.

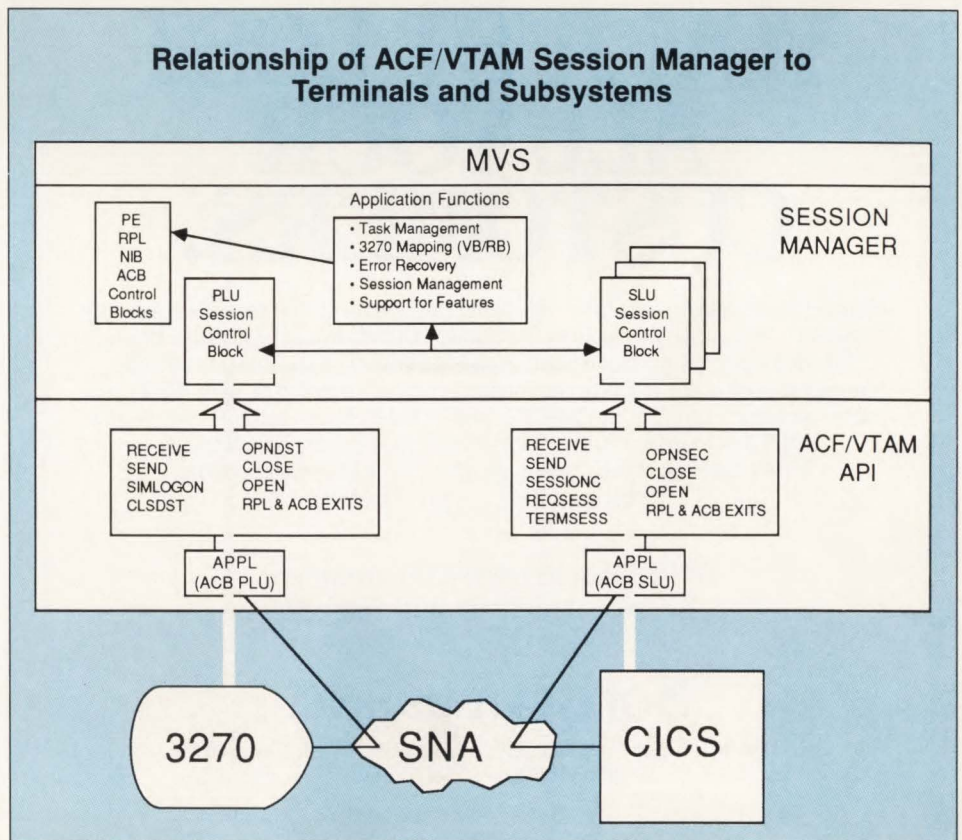
Autolog

The ability to automatically log a user on to applications simplifies the common access problems that users face. Some session managers provide a CLIST language that enables the systems programmer to build autolog profiles. Sessions are established automatically when the user logs on to the session manager. Depending on how the autolog profile is set up, the user is either taken to an application session or to a session menu that lists all the active sessions. In either case, several sessions have been established for the user. The autolog feature makes accessing several subsystems an automatic process saving users' time.

Terminal Lock

Terminal lock and unlock is another security related feature that allows users to soft-lock their sessions while they are away for a meeting or for lunch. The user issues a lock command with a password. The active sessions are not available for use until the unlock command is issued using the same password. This feature is nice when you are running a critical on-

FIGURE 3



line job and you are called away. You do not like leaving your system identification logged on or the job running unsecured. With this feature, you can lock your sessions and leave the job running confident that no one will be able to access on-line information or be able to interfere with your work.

Access Control and Network Image

The last security related feature is access control. Users are provided with

menus to access subsystems and on-line applications and these menus are associated with a user's specific profile. If an application is not on the user's menu, then he will not be able to access that application. Each user's access is controlled by including or excluding items on his menu.

In addition, the session manager can automatically acquire terminals and always retain a session with the terminal. This prevents a user from returning to native ACF/VTAM control where subsystem access is more freely managed.

Menus allow you to present an image of your network and the applications running in that network. Users are associated with menus through their profiles and not the physical terminals. So no matter where a user logs on and gains access to your network, he sees a consistent and familiar menu. The users do not need to know the network or network commands to access applications.

This image remains intact when you must change something in your network. If, for example, you must move a subsystem from one location to another, you can do so without changing the image presented to the end user. Session managers' menuing systems allow you to work on your network without disturbing its appearance. ACF/VTAM USS tables lack the flexibility and adaptability needed to maintain a consistent network image and to provide end users with easy access to applications.

Help Desk Features

One area where the software solution for session management really outshines the hardware solution is in the help desk features. These include peeking and watching other users' sessions, giving and taking sessions, remote input, scrolling, sending messages to individual users or groups of users and many others.

Peek and Watch, Give and Take and Other Features

Peek, watch, give, take, remote input and scrolling are features that allow systems personnel to help end users with access and application problems. You can also use combinations of these features in many other situations that are not directly related to problem solving.

For example, you could conduct classes from a central terminal and have your screens displayed on a number of student terminals. These same features could be used in meetings where information could be echoed to several terminals and viewed by the meeting attendees.

Watch allows you to observe the inputs and outputs to another terminal. Combined with remote input, you could step a user through a problem and correct his mistakes in real-time. Peek is similar to watch, except that peek only provides a snapshot of the user's current screen. When you watch, your screen is updated each time the other user's terminal image changes.

Give and take allow you to take a user's application session to correct a problem and return the session to him after you

See Manager page 46



"POOF!"

YOU JUST LOST ALL YOUR CUSTOMERS.

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

JPU/E Data Recovery Environment.
Because Accidents Will Happen.

SOFTSYSTEMS, INC.

311 Mallick Tower • One Summit Avenue • Fort Worth, Texas 76102
800-331-7846 • 817-877-5070

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

IBM mainframe application software has a new champion

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

Lawson: Leader in software technology.

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

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

Solid application technology.


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

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

Service and support — cutting the "big guys" off at the knees.

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

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



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

Lawson, IBM mainframe software's new champion.

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

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

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

Accounting

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

Human Resources

- Payroll • Personnel • Benefits

Distribution/Procurement Management

- Order Entry • Inventory Control
- Purchase Order

LAWSON

...WHATEVER IT TAKES!

For more information, call
Sue Weinacht
(612) 379-0258

LAWSON ASSOCIATES INC.
1300 Godward Street, Minneapolis, MN 55413

CIRCLE #32 on Reader Service Card

MVS Cross Memory Services – Handle With Care

By Bill Carico

Author's note: The terms cross memory and XMS will be used interchangeably throughout this article and those articles following to refer to Cross Memory Services. The acronym CMS is intentionally being avoided, so it will not be confused with VM's Conversational Monitor System.

Micky and I were enjoying our trip overseas and our technical lectures were in their sixth day and going extremely well. As we set the storage dump and the technical manual that we brought with us down on the table in our hotel room, I looked over at Micky and smiled with great anticipation as to what I expected him to teach me. He is a good friend and a top systems programmer at a large data processing center which at that time had just converted to MVS/SP Version 1 Release 3. I was anxious to understand the ins and outs of how Cross Memory Services (XMS) had been implemented in MVS because we were scheduled to give a presentation on XMS to some System Engineers (SEs) the next day during our lectures. Naturally, with so much to see and do in a foreign country we found ourselves waiting until the last minute to prepare for our talk.

I already had a high-level understanding of how XMS had been designed based primarily on some reading I had done. Now it was time to dig into the internals and understand the low-level workings of XMS, so we could prepare our talk and later use the information to update our MVS internals course materials. Although we really did not have to understand the bits and bytes of how it worked, we both wanted to learn as much as we could, so we could answer any questions that might come up during the class.

Micky had already used XMS at his site to upgrade some system modifications and make them compatible under MVS/SP 1.3. These changes had been necessitated by the fact that within MVS, the operator communications programs and message buffers were moved from common system areas and relocated in a new system address space called the CONSOLE address space. In order to scan the operator message buffers, XMS had to be used to access CONSOLE's private area.

Having Micky along on this trip made me think that learning XMS internal design was going to be a cinch. I already had some ideas on how it had been implemented, so I was confident that this was going to be a brief study session.

We both sat down at the table and I began our work session by saying, "Okay, Micky, tell me how it really works and we will be all set for our talk tomorrow."

"I don't know how it works," he shrugged. "I thought you did! You're supposed to be the MVS internals hot-shot."

I tried to hide my newly developed anxiety attack and probed



areas. (See Figure 2.)

3) Data sharing where a program running in a system common area can reference data located in two different address spaces by alternating between them. (See Figure 3.)

So how complicated is XMS? In answering this question, please refer back to IBM's *OS/VS2 System Programming Library: Supervisor manual for MVS/SP 1.3* that contains the following text in the section that describes cross memory under the ominous heading "Warnings and Restrictions:"

"The design and implementation of programs using synchronous cross memory communication is extremely complex. System services use cross memory on a user's behalf; the user can obtain the benefits of cross memory without having to know the details. Using XMS improperly could cause severe system problems. Therefore, it is very important to consider all the implications of using cross memory. Some general considerations that apply to users of cross memory are:

- Real storage requirements might increase
- Resource management is different
- Accounting methods might be affected."

This warning could be paraphrased by saying two things. One, if you are going to fool around with cross memory, you better know what you are doing. Two, as an MVS user you can get the warm fuzzy feeling about cross memory from knowing that MVS developers went to all the trouble to use it so you can benefit indirectly.

What are the Benefits?

XMS provides several benefits: VSCR, efficiency and security.

VSCR stands for Virtual Storage Constraint Relief. This buzzword signifies IBM's commitment to prolong the life of MVS by rearranging the system code and data areas. XMS provides *horizontal* VSCR because it allows programs and data to be moved horizontally from common system areas into private address spaces. MVS/XA, on the other hand, provides *vertical* VSCR because it raises the virtual addressing ceiling of the system. By using XMS, operating system control blocks and programs are relocated from MVS common areas to individually-owned private areas which decreases the system's requirements for common area storage.

FIGURE 2

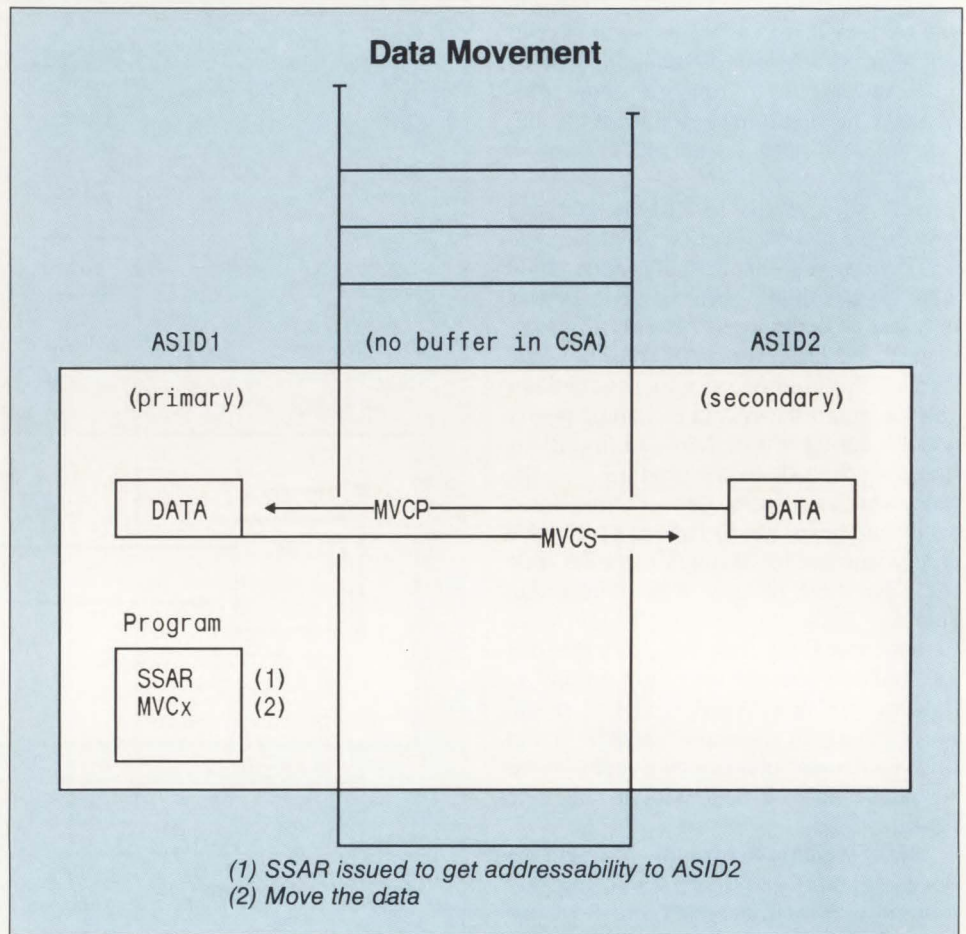
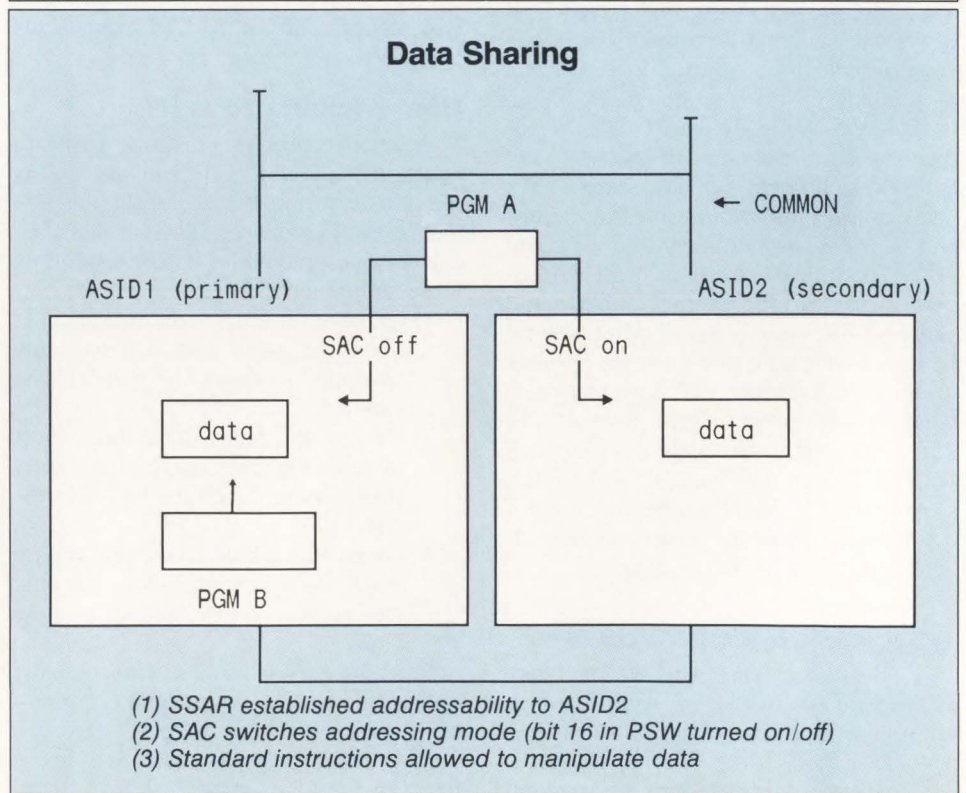


FIGURE 3



Within an MVS address space, since the private area is allocated after the common areas have been assigned, decreasing common-area requirements increases the amount of system private area that is allocated for each MVS address space. Now that XMS is available, new releases of MVS usually contain one or more new system address spaces for MVS components that previously had been located in the nucleus, Pageable Link Pack Area (PLPA) and/or System Queue Area (SQA). For example, up through MVS/XA Release 2.2, the following system address spaces have appeared over the last several releases of MVS since XMS was introduced: PC/AUTH (the XMS supervisory address space), Global Resource Serialization (the ENQ, DEQ, RESERVE functions), TRACE, Allocation, CONSOLE, Linklist Look-Aside (LLA), SMF, CATALOG and DUMP Services.

Second, it allows communication between address spaces to be conducted more efficiently than before. For example, passing data between two address spaces (an activity performed many thousand times a day within your MVS system) without using XMS is a multi-step operation:

- 1) The data is moved from the sender's private area into Common Service Area (CSA)
- 2) The sending program notifies the receiving program that it is delivering some data by scheduling an SRB to post the recipient
- 3) The sending address space typically waits for the recipient to execute and pick up the data being passed
- 4) The recipient, once dispatched and once the data is received, schedules an SRB to notify the sender that the data has been received
- 5) The sender is dispatched and continues execution.

As you can see, this is quite a lengthy process but it was the only mechanism available within MVS for inter-address space communications prior to XMS. Not surprisingly, when XMS was first announced it was reported to use 80 percent to 90 percent less CPU than the SRB/CSA method for passing data between address spaces. This is because XMS provides special instructions to pass information directly from one private area into another, thus bypassing the intermediate step of moving the data into CSA. Furthermore, moving data directly removes the necessity of entering the dispatcher four or five times to dispatch the SRBs and synchronize the programs.

CROSS MEMORY SERVICES TERMINOLOGY QUIZ

Directions: match the definition to the correct term.

Terms:

- 1) _____ Cross memory environment
- 2) _____ Home address space
- 3) _____ Primary address space
- 4) _____ Secondary address space
- 5) _____ Current address space
- 6) _____ Primary addressing mode
- 7) _____ Secondary addressing mode
- 8) _____ Home mode
- 9) _____ Cross memory mode
- 10) _____ PC instruction (program call)
- 11) _____ PT instruction (program transfer)
- 12) _____ SSAR instruction (set secondary ASID)
- 13) _____ MVCP instruction (move characters to primary)
- 14) _____ MVCS instruction (move characters to secondary)
- 15) _____ SAC instruction (set address space control)

Definitions:

- A. Sets the secondary address space to any desired address space.
- B. The mode that exists when at least one of the following is true:
 - The home address space is not the primary address space
 - The home address space is not the secondary address space
 - Secondary addressing mode is active.
- C. Data is fetched from the secondary address space. Instructions might be fetched from either the primary or secondary address space.
- D. Address space where cross memory mode originates. It contains the address space local control blocks that describe a unit of work to MVS. On initial dispatch of a unit of work, the home address space and the primary address space are the same.
- E. A unit of work is in primary addressing mode and HASID and PASID are the same.
- F. The address space whose segment table is used to access data and instructions in primary addressing mode.
- G. Instructions and data are fetched from the primary address space.
- H. The primary address space when in primary addressing mode and the secondary address space when in secondary addressing mode.
- I. The environment in which synchronous inter-address space communication can take place.
- J. The address space in which the segment table is used to access data in secondary addressing mode. In secondary addressing mode, instructions must be in common storage because they might be fetched from either the primary or secondary address space.
- K. Causes another program to get control. The program can be located in another address space.
- L. Issued by a program to explicitly set either the primary or secondary addressing mode.
- M. Returns control from the program called by the PC instruction to the calling program.
- N. Moves data from the secondary address space to the primary address space.
- O. Moves data from the primary address space to the secondary address space.

Quiz Answers: 1-1 2-D 3-F 4-J 5-H 6-G 7-C 8-E 9-B 10-K 11-M 12-A 13-N 14-O 15-L

See Cross page 75

SNA Network Management Using NetView and NetView/PC

By Jay Ranade and George C. Sackett

A comprehensive network management system is needed as the focal point to manage and control all the variables that comprise today's complex networks. NetView is IBM's strategic tool for managing these highly complex SNA networks.

To support a multi-vendor environment, non-IBM equipment and non-SNA resources must be included for true end-to-end centralized network management. NetView/PC provides this functionality under ONM. NetView/PC allows a user or vendor to create an application that executes on a PC and can collect and generate architected network management alerts for non-IBM and non-SNA resources. These alerts can then be processed by NetView or a similar host-based Communication Network Management (CNM) application.

Part I of this article that appeared in the May/June 1988 issue introduced architecture and message flows utilized by NetView Release 1 and NetView/PC Version 1.

NetView Release 2

The scope of NetView was expanded under Release 2 to include management of systems and network operations. In addition, the focus of NetView Release 2 is on distributed control with central management, support for monitoring SNA Node T2.1, functional control of non-SNA resources and further support of IBM's Open Network Management (ONM) architecture by providing an extensive set of generic alarms to eliminate vendor-specific description and recommended action code on the host. (See Figure 1.)

Automation and Central Management

Network automation under MVS is accomplished by the use of the MVS Subsystem Interface (SSI). Executing as a subsystem of MVS, all system messages can be examined for automation. Any NetView console can act as an MVS operator console for issuing MVS system and subsystem commands directly or for executing a NetView CLIST that contains the system command. Likewise, any NetView command or CLIST can be entered on any MVS system console excluding full screen function commands for Session Monitor, Status Monitor and Hardware Monitor. These commands still require a NetView operator's console. Under the VM operating systems, similar functionality of the MVS SSI is provided by the VM Programmable Operator facility (PROP). For the VSE operating system, NetView utilizes the facilities of the program product VSE/Operator Communication Control Facility (OCCF). However, NetView command entry on the VM or VSE system console is not supported. NetView's new position not only includes enhanced network management, but also the management of all computer operations.

Automation Enhancements

NetView R2 automation facilities can be enabled at both the distributed system (that is, IBM 9370) and at the focal point. A major enhancement to the Hardware Monitor is the ability to automate responses to solicited and unsolicited alert notification messages. A NetView

CLIST can be executed automatically based upon additional information in the alert notification message format. This ability combined with the central management enhancements discussed in the next paragraph is a formidable base for future automated control using an expert system.

Central Management Enhancements

The introduction of the IBM 9370 distributed processor is a direct influence on CNM enhancements for NetView R2. The 9370 processor is being touted as an inexpensive, powerful mainframe processor that can be operated by technically untrained personnel. However, central management and control of these distributed processors and other large mainframe processors in the SNA network can be accomplished through several new CNM enhancements. These are the Message and Alert Notification Routing facility, the Message Driven Alert facility, the support for LAN and the Communications Network Management (CNM) Router Function. Each of these functions, either independently or in combination, greatly improves both central network management and automation of a composite network.

Message and Alert Notification Routing Facility

Message and alert information pertinent to a distributed processor can be routed to the NetView focal point operator console by the Message and Alert No-

“Now, my dumb terminal thinks it’s a fax machine!”

VM FAX. No other software package lets your VM terminal do so much. Now save money and time by sending your faxes through your dumb terminals. You’ll feel like you have a fax machine right on your desk.

VM FAX is so easy to use. And, it’s fully integrated with PROFS. Just indicate your recipient by “nickname”—no long string of numbers to remember. VM FAX automatically sends and logs the message. Best of all, the user interface is totally transparent.

Of course, we also let your VM system interact with other communications services. E-Mail. Subscriber databases. Telex. For a free, 30-day evaluation copy of VM FAX, call us today or mail in the coupon.



**SYSTEMS &
TELECOMS**

Suite 300, 12030 Sunrise Valley Drive
Reston, Virginia 22091-3409
Phone: (703) 391-2712
Telex: 249803 • Fax: (703) 476-2217

CIRCLE #151 on Reader Service Card

VM FAX is yours FREE for 30 days.

There’s no modification to VM or PROFS. See how smart your dumb terminal can be.

- YES, rush me at no cost an evaluation copy of VM FAX!
- Please send me further information on VM FAX.
- I’d like to have a representative contact me.

Name _____

Title _____

Organization _____

Address _____

City _____ State _____ Zip _____

Phone _____

Best time to call _____

Send to: Systems & Telecoms Inc.
Suite 300
12030 Sunrise Valley Drive
Reston, Virginia 22091-3409



MJ-788

tification Routing facility. This facility consists of a supplied CLIST that can be modified to route specific messages and alerts from the distributed processor to the NetView focal point. This means, of course, that NetView R2 must also be active on the distributed processor. Based on the information of the message or alert received from the distributed processor, the focal point can execute a CLIST to automate the operator's action for that particular message or alert.

Message-Driven Alert Facility

Any system, subsystem or network operator console message that is viewed by the NetView command facility can now generate an alert to the NetView Hardware Monitor. The Hardware Monitor then treats the alert like any other alert it receives. The alert is displayed on the Alerts Dynamic screen for viewing by the operator. This incorporates not only network hardware and facility fault alerts, but also the operating system and subsystem faults. By using the Hardware Monitor Alert Automation facility, system and subsystem faults can be corrected automatically.

LAN Support

Faults detected on a LAN can now be forwarded to the NetView focal point using the above-mentioned methods. In addition, the fault data has been enhanced to include fault domain alerts, LAN errors occurring on the adapter addresses and a filtering scheme based on the adapter addresses. The extended support for LANs is increasingly important for LANs positioned off a distributed processor.

Communications Network Management Router Function

In VTAM, the CNM Routing table is used to direct the unsolicited network services alerts (NMVTs, RECFMs, RECMSs) to the unsolicited command processor responsible for handling unsolicited alerts. VTAM requires that only one unsolicited command processor can be active for any given network services RU at one time. The CNM Router Function allows the user to send the unsolicited alert to a user-specified component based on the identification information found in the record. This means that more than one unsolicited command processor can be active and receive unsolicited NMVTs from the network.

Session Monitor SNA Node Type 2.1 Support

SNA Node Type 2.1 supports IBM's

Low Entry Networking (LEN) architecture for peer-to-peer communication. This new SNA node type is supported by the NetView R2 Session Monitor in conjunction with VTAM V3R2 only and NCP V4R3 or NCP V5R2. This allows the monitoring of Primary Logical Units (PLUs) that are not under VTAM control over an SNA backbone network. Support of SNA Node Type 2.1 is furnished by new Session Awareness data (SAW) records under VTAM V3R2. The new SAW data contain a unique Procedure Correlated Identifier (PCID) value. The PCID is used to identify each session the PLU

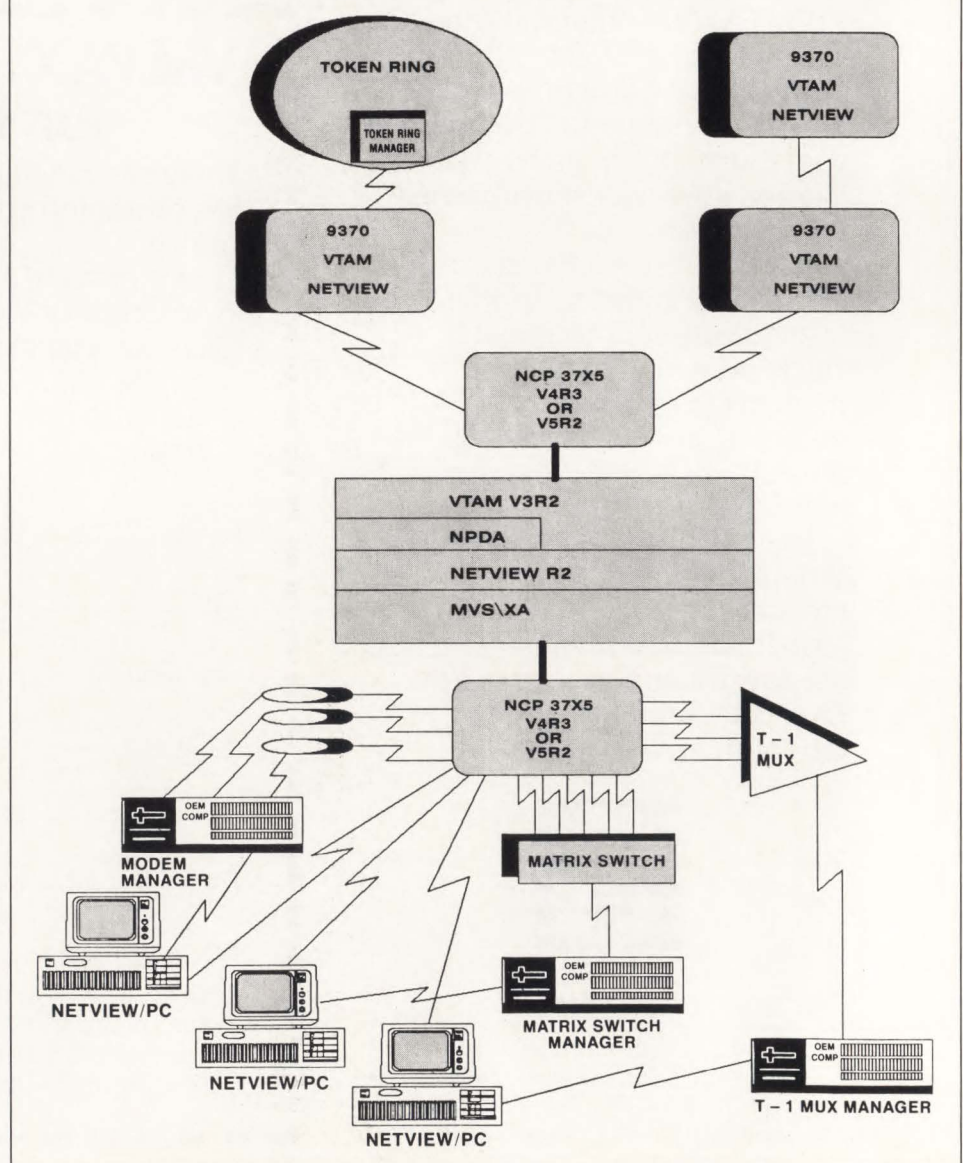
has with a session partner. In this way, independent LUs (with multiple concurrent sessions) of SNA Node Type 2.1 can be monitored and tracked for each session by VTAM and NetView's Session Monitor.

Generic Alerts

Under NetView R2, the Hardware Monitor display panels and messages have been enhanced by the new Generic Alert format of the Network Management Vector Transport (NMVT). The implementation of this new format eliminates the customization of NetView Hardware

FIGURE 1

Network configuration utilizing ONM with NetView R2 and NetView/PC V1R1.



Monitor display panels and messages for product-specific alerts. Instead, the alerts dynamic message, event detail and recommended action screens are built upon the text identified by the generic alert code points. The code points are used as indices into tables that reside at the host on disk. This provides for individualized alert text for each type of alert received versus NetView R1 pre-defined alert messages and displays that were indexed by a components identifier and an alert descriptor code.

In keeping with IBM's Open Network Management architecture, NetView R2 provides a facility for vendors to add their own generic alert code points. A code-point text table can be created to provide product-specific code-point text that can be used with IBM's code-point text to create the Hardware Monitor alert messages and screen displays. Prior to NetView R2, only IBM's text or the product-specific text could be displayed. Using generic alerts, product-specific and IBM code-point text can be combined on the same message and/or screen display.

NetView R2 will, however, continue to support non-generic alert formats that were supported under NetView R1 and the individual CNM program products.

Service Point Command Service

Management and control of non-IBM equipment from NetView is now possible from the new added function of the NetView Command Facility. This new function is called Service Point Command Service (SPCS). Using the command facility, a NetView operator can enter four supported commands that are to be executed by an application program of a service point. IBM's NetView/PC is an example of a service point. The supported commands are : RUNCMD, LINKDATA, LINKTEST and LINKPD.

SPCS RUNCMD Command

The RUNCMD is a generic command that is recognized by the NetView command facility and is transported to the service point for processing. For example, RUNCMD can be used by the network operator to obtain the status of the vendor's equipment. Command syntax checking is not performed by NetView and must be in agreement with the vendor's application executing on NetView/PC. The vendor's application must be coded to handle command errors, as well as execute the requested command. The

response from the service point is displayed under the command facility function of NetView.

SPCS LINKDATA Command

This command is used to request control, error and statistical data maintained by the service-point application. The service-point application responds to this command with an architected record. The information supplied in this record is displayed by NetView in full-screen mode.

SPCS LINKTEST Command

Using this command from NetView R2, an application executing on the service point will request DCE tests. The results of the tests will then be sent to NetView in an architected record for presentation by NetView in full-screen mode.

SPCS LINKPD Command

Problem determination analysis can be performed by the service-point application for a given element. Using the LINKPD command, the network operator can request the service-point application to perform alert correlation for a specific problem. The response to this command is a generic alert. The generic alert is then processed by NetView R2 in the same manner as solicited and unsolicited NMVTs. In this case, the NMVT is solicited.

NetView/PC Version 1.1

Support for the enhanced alert and command extensions of NetView R2 is provided by the IBM service-point product named NetView/PC Version 1.1. This new release of NetView/PC enables NetView to extend network management and control to non-IBM equipment. The intent of this release is to further solidify IBM's Central Network Management architecture by having NetView/PC support both generic alerts and SPCS.

Generic Alert Support

NetView's new generic alert architecture places more responsibility on the vendor service-point application than the previous non-generic alert format. The current number of code points supported by NetView R2 without vendor-specific code-point values is quite extensive. Each piece of text displayed on the Hardware Monitor screens can be selected by code-point values. This increases the detail of code needed in the vendor application to support all the possibilities. The current base plus further expansion of the code-point values to increase the detail of the alert message may become overwhelming.

Service Point Command Service Support in NetView/PC V1.1

The previous release of NetView/PC was used by most vendors to consolidate their DCE alarms on NetView. With this release, a NetView operator can in fact control the DCE equipment from the NetView console using SPCS. The vendor's application must be developed with the new API/cS (Application Program Interface/Communication Service) interface to support the encoded commands received from NetView. Partial or perhaps total operational and managerial control of their equipment by NetView is possible. Of course, this support is dependent upon the vendor's willingness to provide this service to the customer. By providing the capability of NetView to actually operate and test non-IBM DCE by using NetView/PC, IBM is supplying the means for total network management.

Token-Ring Network Manager for NetView/PC

The IBM Token-Ring Network Manager V1.1 for NetView/PC V1.1 executes as an application under NetView/PC. Alerts generated on the token-ring network can be automatically forwarded to the SNA focal point application (that is, NetView). The alerts are of the generic alert format. The alert forwarding from a token-ring manager executing under NetView/PC coupled with SPCS functionality can therefore provide work group management from a NetView console. ☹

ABOUT THE AUTHORS

Jay Ranade is assistant vice president in the Advanced Office Systems Group at Merrill Lynch and Co., New York. He is also editor-in-chief for a new series of books on IBM and DEC hardware/software for McGraw-Hill Book Co. He is author of two best sellers, VSAM: Concepts, Programming and Design and VSAM: Performance, Design and Fine Tuning.

George C. Sackett is a leading communications consultant and the president of ASAP Technologies, Inc. that specializes in providing consulting and educational services and is located at 297 Orient Way, Rutherford, NJ 07070, (201) 939-6821.

Ranade and Sackett are writing a two volume book titled SNA Networking: A Guide for Using VTAM/NCP scheduled to be published by McGraw-Hill in September, 1988.

Software Maintenance: Undermanaged & Understaffed

By Jan Snyders

Software maintenance is gobbling MIS's resources. According to the Quality Assurance Institute in Orlando, FL the defect rate in software maintenance is three-to-ten times as great as new system development.

The conclusion of the Institute's 1985 survey showed that in two-thirds of the organizations surveyed, no one was responsible for managing the software maintenance process. Another problem was understaffing.

Some of the survey's conclusions were:

- Although maintenance accounted for more than half of the data processing budget, approximately one-third of the organizations had an individual assigned the responsibility for the software maintenance process.
- Of seventy-eight percent of the organizations surveyed, each organization had to have systems maintained by specific individuals because no one else in the organization understood the logic of the software programs.
- If the requests for maintenance stopped, it would still take an average of 22.8 months to complete the existing maintenance backlog and

between two and 60 months to complete the backlog.

- Approximately two-thirds of the organizations surveyed had a software maintenance development methodology. Most of those organizations obtained their methodology as part of their new system design methodology.

Ironically, the problem has not improved in the past three years.

To assist readers in learning more about software maintenance, MAINFRAME JOURNAL asked nine industry leaders for their opinions on the problems and possible solutions. Included in the input are: Charlie Bachman, president and chief executive officer of Bachman Information Systems (Cambridge, MA); Paul Bassett, vice president of research, Netron (Toronto, Canada); Gerald DePass, president, Adesse Corp. (Danbury, CT) Howard Feinman, director, software products division of Catalyst, an information technology firm of Peat, Marwick & Main, (Chicago, IL); Phil Loiterstein, president, Group Operations, (Washington, DC); Mario Morino, chief executive officer, Morino Associates, (Vienna, VA); Richard Nashleanas, president, Eden Systems

(Carmel, IN); Kent Petzold, president and CEO, Viasoft (Phoenix, AZ); and Marc Sokol, executive vice president, Realia, (Chicago, IL).

MAINFRAME JOURNAL: Is maintenance still a problem and how much?

Loiterstein: The maintenance problem, in fact, is increasing rather than decreasing. The days of thinking of maintenance as fixing bugs are long past.

Bassett: Maintenance is any activity performed on software after that software has been put into production. Basically, maintenance comes from the changes in business. Every time a new system is developed, the maintenance task is increased.

Loiterstein: At one time we could take a year to develop a system, hang around for three or four years and something else would come in. That cycle has been extended. Now we still take a year to develop it, but suddenly it's staying around for six or seven years. In this area, it really is 50 to 90 percent of the programming budget.

Sokol: Up to 80 percent of all mainframe resources are consumed in the

maintenance area and only 20 percent for new application development. Because so many resources are put into that area, making those resources more efficient and more productive is very much of a problem and a goal for data processing and programming managers.

Nashleanas: Software maintenance is a hidden problem. Management hides it in various overhead costs such as staff time, computer time and so on. Management doesn't feel the cost of maintenance as taking money out of its pocket to actually pay for maintenance every time it's needed. Management knows it's spending a lot for maintenance and it just doesn't know quite what to do about it.

MJ: *There are claims who from 20 to 50 percent savings can be realized with some of the new tools and methods. This is one of the most heavily debated topics.*

Some of the cost cutting methods available include assigning a person or persons to the maintenance task, using tools such as 4GLs, COBOL unscramblers or change control packages. One of the oldest attempts to reduce the software maintenance cost was having a person or persons, depending on the size of the installation, devote all of their time doing maintenance whether it be upgrading or simply fixing bugs. Is this fact or fiction?

Bassett: Maintenance, psychologically, is a very depressing task. Programs were never written explicitly to have maintenance easy to do. It's a difficult task relative to building something with no constraints. That is, you just invent something from first principles. You're doing maintenance and you have to somehow fit in with the surrounding programs. That puts a heavy constraint on you. The detective work a programmer has to go through to locate all the different fragments that are relevant to the changes is a tricky and a frustrating task because things are missed. The only way to find out is through a torturous process of trying it, seeing what went wrong and trying it again — almost by trial and error. That can burn people out quickly. So having people do nothing but maintenance day in and day out is emerging them in a psychologically negative environment.

Loiterstein: That's right. It is not a good idea to assign specific people to the job of maintenance. Everything the developers developed was always bad and they never learned what was wrong or what didn't work. Even when the developer put out new models of programs that had to be maintained for a year or two,

they still make the same mistakes. There isn't a lot of communication between the developer and the person maintaining it. It's like the developer became the higher-level person and maintenance is given to the less talented person.


MJ: *There are, however, those who think the idea of maintenance specialists is beneficial to the company.*

DePass: I think if the job is described properly, the talent is paid for properly and the right person is put on it, it's the right thing to do. To scatter maintenance around and treat it as some sort of step-child is a mistake. Maintaining a system is an important activity.

Feinman: Yes. It is critical to put maintenance specialists in place because you want people who really understand the programs they are looking at. Programs are not consistent, since there are tricks done in terms of coding and programming that are not standard within programs. So, organizations put people who are specialists on maintenance. These specialists become familiar with the programs because they have to understand what's happening.

Sokol: The specialist is more in tune with where the industry is heading regarding automated tools and methodology. As an example, if you have a huge project that you are about to undertake, such as to take a program written 15 years ago and enhance it in some particular manner, you could consult with the maintenance specialist and ask him what he thinks is the best way to go about it. You could ask about COBOL unscramblers or conversion to higher-level languages. That kind of resource would be really beneficial.


Nashleanas: There are a lot of advantages to having maintenance specialists. They're technically higher quality. They're better at doing maintenance since they know how to do it, and most importantly, they have a longer-term outlook on doing their job which is really what you need. However, on the negative side, maintenance specialists tend to ask questions. What is my career track? Where do I go from here? Is this valued by the company? Does it have the same type of exposure to get ahead and get me where I want to be? There is the opportunity to fall into the indispensable-man syndrome. That is, you're so good at doing your job that you can't afford to get promoted or be moved anywhere else because you're the only one who's good enough to maintain those crazy things.




Paul Bassett



Charlie Bachman



Gerald DePass



Phil Loiterstein



Mario Morino



Richard Nashleanas

Petzold: There are really two predominant organizational structures. One structure is dividing the shop between new development and maintenance. The maintenance side grows larger and larger because this group has lacked technological solutions. An advantage of this structure is that it provides for focused product training and focused cultural development of the two groups. The other is application responsibility. This environment is popular because it makes sense to put the best resource on a given project. Best in this case, meaning the people with the most application experience. In this instance, managers have been forced to almost hold captive particular programmers because of their knowledge and understanding of the applications.

Morino: An unusual approach to this is found in a service bureau in Canada that is a consortium of services to charter members. Three large firms in Canada use this center. They built dedicated support teams for each company that they service. There is a lot to be gained in this approach because now you have people who are working with the systems on a continuing basis, so they're not going to learn in curves. They tend to understand the systems better and they can build better relationships with the end users. You end up with a much more effective process.

MJ: *Are 4GLs being used as an aid in solving the maintenance problem?*

DePass: I believe it's a fact that higher-level languages and tools that work with them simplify the maintenance process. It also tends to spread the work out among a lot of people because the users do their own thing, so maintenance, per se, falls in their laps.

Feinman: The 4GLs are easier to maintain because they work, in many cases, on the assumption that you maintain it at a spec level. You don't change COBOL code, but generally drop it. With the 4GL you go back and regenerate a new program that had the old function that you had just coded in that language and add the enhancements. In that case, they are easier to maintain but the big inventory that is running on a day-to-day basis is not addressed.

Morino: Yes, systems written as 4GL applications do ease the problem because they are going after the core part of the problem — systems that don't have enough flexibility and aren't done in a generalized enough fashion — and that's what creates most of the maintenance issues today. The trouble is that you're only

going to alleviate the work on those things that are written in that application.

Petzold: I think 4GLs were a panacea of yesterday. They really haven't met the challenge. They are typically relegated to the quick-and-dirty applications, whereas COBOL is still the king even after 30 years of existence.

Bassett: I don't think 4GLs have made a significant dent in the maintenance problem. First, they are aimed at end-user computing and are really not suitable for industrial-strength, high-quality software systems. They lack the fine tuning abilities that quality software demands in which you have subtle trade-offs among different design constraints, optimization problems and other highly technical spit-and-polish to put into a system. The 4GLs don't lend themselves to that type of fine tuning. The second reason is that 4GL is a programming language just like COBOL and to that extent, like any programming language, supports evolutionary change. Unfortunately, 4GLs along with COBOL and other languages have been designed with the idea that software is a perfectable object. The only thing you do is fix bugs and put it into a restorative state; so, the languages simply have never been designed to cope with that ongoing need to keep changing software in ways that are unforeseen. As a result, when you have to make changes, they're done in an ad hoc fashion. When they are doing ad hoc, again and again, that ad hoc builds up in the code until finally it becomes an impossible chore to sort out further changes.

Bachman: My impression is that it's helping but it's not helping enough that the total resources are being dedicated and still growing. They're slowing down the tide of the maintenance percentage of total resources, but it has not stemmed it in any sense. Most programs written in 4GL are being used for special report preparation and, in some sense, it's being done by people outside of the data processing department.

Sokol: Because these languages tend to be high level and therefore need fewer lines of code than COBOL would, you immediately think they should have less maintenance. However, one thing to keep in mind is that you need to have experts in those particular languages. Whereas there are many COBOL programmers and most of them are experts, there aren't too many experts in those particular languages. Therefore, it may be harder for large companies to find the people to do that maintenance.

Nashleanas: If the 4GLs are used the way they were designed to be used, they have significant benefits in maintenance. At the same time, if you develop an unmaintainable piece of 4GL code, try to pack too much in a piece of code and if you go ahead and generate that, it is as complicated as the original piece. You can go ahead and design in all the same problems that you did with the 3GL.

MJ: *The software maintenance aids, COBOL unscramblers and change control packages are relatively new in the marketplace. How much can they help in reducing the maintenance problem?*

Sokol: When people didn't know about structured programming and didn't understand the psychology of understanding programs, programs were messy and a lot of verbs were used that later turned out to be difficult to maintain. Techniques were used that people frown upon because they're machine specific or compiler specific. For these reasons, I think that the COBOL unscramblers are really beneficial. I think they're the first step to really automating the maintenance process. Descrambling programs by hand is a time-consuming task. An additional benefit is that you can use lower quality level maintenance programmers to maintain even complex applications.

Feinman: COBOL unscramblers help reduce the maintenance effort because they make everything consistent within an organization in terms of the code's structure. Again, by using a restructuring tool and getting that consistent, structured program, you can put programmers on almost any program and they will feel comfortable with what they're seeing and what they're going to see across the system and all applications.

Loiterstein: I think unscramblers help in maintenance but they are no panacea. Today the problem is not the number of lines of code, but it is analyzing the old system. The programmer spends time analyzing the old system which is usually very large and the original developers are not around. With the COBOL unscramblers there is a consistent architecture so the programmer doesn't have to look at the program or the architecture of the program to see how its done. Now it becomes a canned program that does things the same way every time.

Nashleanas: While these tools can yield benefits in maintenance, you have to consider the cost. You have to selectively apply those types of tools to only the code

that truly needs the reformatting. Often these tools come with the analyzer that says whether the program should be restructured or not. If you restructure a well-structured piece of code, you'll probably end up with a piece of code that's not quite as good.

Bachman: Unscramblers help some, but one of the real problems we've found with them is the way they take the comments that have been inserted in the program before. They are less able to put those comments into a reasonable position; so, some of the programs have been restructured but the comments are a disaster because they aren't where you'd expect to find them. Therefore, they don't always achieve the kind of help in maintenance they are supposed to. That is, they do technically what they're supposed to do, but they don't always help maintenance that much. Essentially, a human still has to go back to the program and sort through it to figure out what the change was.

MJ: *Is change control easing the maintenance problem?*


DePass: Change control can absolutely improve and dramatically impact the problems created by sloppy maintenance. I think it's critical. The problem is that people short-change this area. Lots of people pay lip service to change control, but don't apply the resources that make it happen or develop the procedures and make sure they stick by them. Change control can help regardless of whether they are using 4GLs or basic Assembly languages.

Sokal: I think change control is initial among itself. It's something that is crucial and it is needed to do problem maintenance and to make maintenance efficient. It really is required when you have more than one person maintaining a single application. As soon as you have multiple people doing it, you have to keep track of who has the latest version of this particular file, or copy record or who made the last change to the program. You really have to structure the revision history of who did what to what. I think automating maintenance is really required.

Bassett: The hardest problem in maintenance is that when you go to make a change you find that the change you want to make is not localized in one part of the program. There are fragments of text that are scattered throughout a thick program listing. The programmer does not know when he starts to make changes how many fragments there are that are relevant to the



Kent Petzold



Marc Sokol

change being made. In addition, he is never sure whether they have found them all. Change control helps with that problem because at least when you do find the pieces that you think are relevant to a change, a change control tracks them for you.

Morino: Although change control helps the integrity of the maintenance process in the sense of doing exactly what it implies, it does nothing to reduce the maintenance workload. You've still got all the work to do to get the changes, the design and so on. Change control puts the integrity piece on the back end to make sure that after you've made these changes, they're going to get into the system properly.

Petzold: Change control, at the pro-

gram logic level and in terms of reusable code where the programmer works and where time is spent, have a long way to go. Change control software is used to assist MIS management in tracking the status of the many maintenance projects in the department and there's been good advances in this area. But little has been done in the way of new technology to make the maintenance programmer's job easier and more productive.

MJ: Is there a solution to the maintenance problem?

Petzold: The solution involves the process of understanding change and the effects of the change on a company's programs. Up until today no single vendor has applied the technology and methodology required to understand all of the capabilities involved in effectively managing all of this change.

Bachman: There is not a solution in the marketplace. Some people are working toward that solution by providing tools to help people solve the maintenance

problem, but there is no complete tool out there.

DePass: I don't think we're going to get rid of the maintenance problem, unless we stop the world and I guess at that time we all get off. The maintenance issue is going to be with us as long as we have programs. The only thing one can do about it is to make an attempt to control it and minimize change.

Loiterstein: Before the maintenance problem can be solved it must be recognized that there is a maintenance problem. We have to say, "There is a problem and this person needs different tools." After many changes are put in hastily, the code deteriorates. I think structuring in products does help maintain the code, but if you don't have automatic documentation tools that can maintain your needs, you're in the same place you were before. There are a lot of tools that the person needs, but right now they're not being developed.

Sokol: Maintenance can be compared with the question, "When will COBOL die?" Until there exists some sort of system that can read a COBOL system, understand what it does and represent it at some high level or super-high level method, there will be problems. COBOL unscramblers do help because you at least need to be able to understand the control and data flow to be able to do some maintenance.

Morino: The answer is better management. It's not tools. Tools never solve a problem. People using the tools solve the problems. An optimizer or a planning model by itself does not solve the problem. A manager has to know what is needed to solve the problem and then make sure that he understands what he bought and brief the users on how it is used.

Nashleanas: The problem can be solved by using today's solutions. If I were a manager, I would ask questions such as what has to be done and what can I do? The first step would be to implement change control in such a way that degraded programs can't be degraded by maintenance. The second step would be to re-engineer and attack the poor-quality programs and get them up to snuff.

For a list of maintenance software vendors and their products, circle number 389 on the Reader Service Card. ☉

JOB SCHEDULING

... the choice is yours!



Classically Aged scheduling systems, written for the 70's.

OR . . .

ASE, designed for the 90's!



- Totally automates job setup including run dates, data string changes, day-of-week and/or month-of-year variables, and optional user specified prompting.
- Automatically tracks all production jobs, manages system resources (tape drives, initiators, critical paths, & more), and dramatically reduces operator intervention.
- Full operator-command support allows any MVS or JES2 command to be scheduled as easily as batch jobs (with all ASP dependency support).
- Uses full ISPF/PDF implementation (split screen, edit/browse, table displays, line commands), with no new editors, functions, or philosophies to learn.
- Extremely efficient (typically less than 1 percent total overhead), nearly automatic 1 hour installation, NO MODS, NO EXITS, NO HOOKS, NO IPL!
- Lowest cost-per-function of any scheduling system available today. Introductory price(s) starting at \$10,000 (with a 30 day free trial).



Chaney Systems Support, Inc.
W18484 Kingston Drive
Muskego, WI 53150
(414) 679-3908

CIRCLE #146 on Reader Service Card

ABOUT THE AUTHOR

Jan Snyders, CSP, is a free-lance writer specializing in software.

Systems Application Architecture (SAA): An Overview

— *By Phil Layland* —

It has been said that Systems Application Architecture (SAA) is the most significant IBM announcement in a decade. It will surely affect our lives in the data processing world over the next decade or more.

At this stage SAA is just a concept. Some pundits believe it will take up to two years before the first SAA products are delivered and until the end of the century before any significant proportion of the product range will be SAA capable.

Should you be concerned with SAA today? Should you be aware of SAA when making future decisions? When products are delivered with SAA capability, will that be the time to start investigating?

Understanding the concepts now will be excellent preparation for future decision-making when SAA products are delivered.

What Is SAA?

Systems Application Architecture is a collection of selected software interfaces, conventions and protocols that will be the framework for the development of consistent applications across future products.

See SAA page 74

SAA is a statement of direction for long-range planners. At some point in the future, their applications will run in any environment encompassed by the architecture. Today's SAA gives the foundation on which to build.

Water On The Rock

A performance management and capacity planning primer

By R. Douglas Swords

This article is not directed to scarred veterans of the performance wars or computer performance scientists. It is directed to individuals with little or no experience in performance management who would like to establish a program and need a starting point.

Mandate:

Reduce and Control Costs

Money is tight, especially if you work in an organization that provides health care. For many years the cost of health care has been rising. This explosive increase in health care is all over the newspapers, the subject of congressional inquiry, the number one concern of elderly citizens and people with chronic health problems, the subject of studies by health insurance organizations and, recently, the subject of restrictions by the federal government. The mandate from the government, insurance companies, the general population, elected officials and hospital administration is simple: reduce and control costs.

There are areas within a health care organization that are less constrained and more variable; usually those areas dealing with direct patient management. The patient will get what he or she needs in terms of treatment. This is a variable and, within limits, is subject to fluctuation. This is the area addressed by government regulation in terms of recognized averages in the treatment of certain illnesses and the subsequent reimbursement of the health care organization by the government for individuals receiving treatment who are covered by Medicare and Medicaid. The bottom line here is if you exceed the average and there are no extenuating circumstances, you may not get reimbursed for the above average cost. This is not my area.

My area is information systems. This area is closely watched by everyone for one simple reason: it is the only area that has a potential impact on *every* other area in the hospital. The cost of system resources is high. For our purposes this cost includes hardware, operating systems and applications. Our charter, as capacity planners and performance managers, is to optimize these resources. We have to squeeze every bit of capacity out of what we have and put off acquisitions of additional resources until the last possible moment. We must be certain that what we are requesting in budget is a legitimate requirement. The question is, how do we do this effectively?

Performance Management and Capacity Planning

The simple answer is establish a performance management and capacity planning program. Someone may wonder why we say "performance management and capacity planning" in the same breath. You cannot capacity plan from an untuned system. To do so would be to build in inefficient use of resources and thus reduce the dollar-to-resource ratio that is critical, especially in health care.

In any untuned system there may be as much as 35 percent unused and available resource. So it is true to say that while system tuning may make some job or transaction run faster, tuning also means the optimization of existing resources to the fullest possible extent. I know this from my own experience in my environment. Out in the rest of the world the number may be higher or lower, but you can bet if there has been no tuning then there is untapped resource available.

In most cases performance management and capacity planning software tools are invaluable in helping to identify where

these inefficiencies might exist. However, your management may hesitate to give budget support for such tools especially if they do not believe there would be a benefit. It is difficult to convince someone, without concrete evidence, that the expenditure of \$80,000 might in fact postpone the acquisition of several hundred thousand dollars worth of hardware upgrades.

There are some steps you can take without this type of software to tune parts of your system and get your proof. These *basic* steps will establish that tuning is essential, improvements can be measured and in-house staff with the appropriate tools and desire can perform these functions.


For the average environment that includes CICS and some amount of batch work there are several basic areas to study and perhaps realize some surprising savings in resource consumption. A batch job consisting of several job steps with multiple sequential files processed is the best candidate to start with.

Dataset Blocking

Dramatic results in improving batch throughput and reducing the elapsed run time of batch jobs can be realized from some work in the basic area of dataset blocking. Many people do not believe it. Even though new DASD is extremely fast as measured in terms of seek times, average rotational delay and full track rotation dataset blocking is still critical.

Let me illustrate with the following example, one that you can recreate in your own environment. I used a COBOL program that reads a tape-based dataset with a record length of 31 and a blocksize of 31000. The output dataset was created on a 3380E first with a record length of 31 and a blocksize of 31. A counter in the

Having one VSE console is fine... if you run one job at a time.



Unfortunately, you don't get to work that way. Instead, you're forced to run multiple jobs from a single VSE console.

Fortunately, there *is* an answer.

Because now you can operate your system *from up to 48 true consoles*—local or remote. With SMARTECH's Display Operator Console Support (DOCS).

And DOCS shares your terminals with your other on-line software—with instant, hot-key switching. DOCS even tells you which jobs are running. Which jobs aren't running. And why.

So using DOCS means your programmers can check on jobs from their own desks. For a more organized computer room.

What's more, users can communicate with VM directly from the DOCS console, so you can manage your

system with fewer terminals. Through the IUCV Interface, you can manage multiple VSE systems under VM—from one console—even in VAE mode.

And DOCS lets you custom-design security parameters for your entire VSE system—without requiring an on-line system such as CICS.

What's more, DOCS practically installs itself in 30 minutes or less—and requires no assemblies or customization. It comes completely ready to serve you.



SMARTECH backs all these features with a commitment to service that's unmatched in the industry. A commitment that supports all releases of VSE, VM and MVT.

Join the more than 400 users who enjoy increased efficiency with DOCS. Call David Shepherd at 1-800-53-SMART for more information.

Because multiple jobs need multiple consoles.

Call 1-800-53-SMART for your free DOCS trial.
Outside the U.S., call (214) 630-TECH.



SMARTECH SYSTEMS, INC.

Turning high technology into *SMART* TECHNOLOGY.SM

Corporate Headquarters: SMARTECH Systems, Inc.
8701 Carpenter Frwy. • Dallas, TX 75247 • Telex: 9102503110

© Copyright 1988, SMARTECH Systems, Inc. VM, VSE and CICS are registered trademarks of the International Business Machines Corp. MVT is a registered trademark of Software Pursuits International. SMARTECH and DOCS are trademarks of SMARTECH Systems, Inc. Copyright © 1988 SMARTECH SYSTEMS INC. All rights reserved.

CIRCLE #65 on Reader Service Card

program was set to stop processing after one million records had been read from the tape and written to DASD.

On the first run after two hours of clock time we got tired of waiting and killed the job. It had consumed more than 800 cylinders of disk and still had not completed running. On the second run the disk dataset was created with a blocksize of 5859 characters or 189 records per block. Run time was 2.27 minutes and it required 51 cylinders of disk.

We then recreated the disk dataset with a record length of 31 and a blocksize of 23467. The job ran in *one minute and 39 seconds* of clock time and consumed *only 45 cylinders of disk*. For the performance scientists waiting to pounce on me, this test was done under controlled conditions. This may seem to be an extreme example but it proves the point: block sizes are critical in reducing elapsed times, volume-busy trips to the channel and disk space requirements. On 3380 DASD the magic blocksize is 23476 or half a track, so aim for that when determining your block size.

If we examine why this is we find that when the dataset was blocked at 31 characters or one record per block, each logical write in the program generated a physical write to the dataset (buffering not considered). So, to process one million records you had one million writes. When the blocking factor was changed to 5859 characters, there were 189 records per block and the ratio was 189 logical writes to one physical write and at 23467 characters or 757 records per block, we had 757 logical writes to one physical write. Thus, to write one million records at 757 records per block we wrote 1321 physical records or blocks. Think about it.

In some cases you can see further improvement by using the BUFNO keyword in the DCB parameter. Experiment with adding buffers starting at 10 then moving up. This will also substantially reduce head movement at the volume level. I have done this and have measured 70 percent reduction in head movement activity on one volume where many sequential datasets reside.

If you are in an environment that has tens of thousands of datasets, it is not practical for one person to try to change them all but you also have the most potential for improvement. Just pick one particular job, average its before blocksize modifications elapsed times and then change the blocks and rerun. Report it to your management. Then ask for money

for a product like BlockMaster from Computer Associates or Sequential I/O Optimizer from System Connections to help you reblock the others. Why? This action may save the purchase of some number of volumes of DASD or at least allow you to place more datasets on a particular volume. You may also be able to complete more batch jobs in the same slice of time.

Manipulating ICS and IPS Members

Another method of improving batch turnaround depends on the physical construction of a particular job stream. For example, if the first job in the stream creates a dataset that the second and third use as input, you could place the first job in a different performance group through the use of the Installation Control Specifications (ICS) member in PARMLIB.

*We have
established
a successful
program here,
made all of the
mistakes and some
of the progress.*

R. Douglas Swords

Using a higher priority performance group could allow the first job to complete faster, further reducing the overall elapsed run time. Here we switch Installation Performance Specifications (IPS) members through the AUTOCMDS facility at midnight to allocate more resources to batch processing when the CICS activity is at its lowest. There may be rules about this in your shop so check around first. Manipulating ICS and IPS members is more complex and it really is not a good starting place for a beginner.

We can also examine Partitioned Datasets (PDS) in a basic sense. If you think of a PDS as a warehouse where many different things are stored and the directory of the PDS as a list of where those items are, you can see that by placing the items in greatest demand nearer the list

means that you travel a shorter distance to get it. Reorganizing a PDS based on the fetch rates of its members is an effective way to improve response times for TSO users and the loading of modules from load libraries for CICS. There are tools to help you do this and nearly everyone who sells DASD performance products will have such a tool. If you have trouble getting something like this you can always break up very large PDSs and place highly fetched members in smaller PDSs together.

CICS Performance Monitors

It is difficult to analyze CICS without some sort of performance monitor such as the Omegamon CICS from Candle, the Monitor for CICS from Landmark or CICS Manager from Boole & Babbage. These firms will offer products and some assistance for a trial period. This is a good way to build a case for the purchase of the product.

There are different management techniques depending on the level of CICS you use and whether or not you run under MVS/XA. The non-XA shops have more to gain by careful management of the storage within the available region size due to Virtual Storage Constraint (VSC).

Within CICS the very basic things to consider are program compression, allocation and utilization of the Local Shared Resource (LSR) buffer pool, evaluation of page size and perhaps evaluation of resident programs.

From the standpoint of the LSR pool, consider that datasets with high update and browse ratios should be removed and this removal may yield better utilization by the high read-to-write ratio datasets. Some of this can be determined from your CICS out list statistics. Also, when allocating your buffer pools you do not want to over or under allocate. Try for something in the range of five percent buffer wait.


If you never see any waiting for buffers then you have too many. Much more than five percent will degrade performance. For any datasets removed from the LSR pool, you will have to do the same evaluation for its individual strings and buffers.

Program Compression

Program compression can significantly affect the performance of CICS. In a basic sense, program compression occurs when a program is needed that is not loaded and there is no room to load it. Programs not currently being used are flushed out and

See Water page 83

At \$1,595 THE MAJOR DIFFERENCE BETWEEN OUR 3270 PRINTER AND THEIRS.




OURS EXISTS.

That's right. Take a minute and flip through this magazine. You won't find another mainframe printer that gives you as many features for anywhere close to this price. Starting at only \$1,595, the Prima-CX is designed for industrial-grade applications, not PC-type printing. It attaches to IBM mainframes as well as 4300 and 9370 mid-ranges — via IBM 3174, 3274, and 3276 cluster controllers.

So whether you're producing memos or listings, there's a Prima-CX model (80- or 136-column) that's for you. The IBM-type front panel puts control at your fingertips. You can feed cut sheets without removing continuous forms. Program tear-off and RPQ options. Cancel print jobs from the printer. Take advantage of the 220 cps draft speed. You'll also enjoy the low noise level — less than 55 dB. All this in the smallest footprint of any coax compatible printer.

No one else offers a first-year, on-site service contract for only 49 cents.

They'd be afraid to. But, buy a Prima-CX professional printer and we'll give you a comprehensive, *first-year, on-site* service contract for *only 49 cents!** We can make this incredible offer because our design ensures maximum reliability. Contact us today for complete details.



Printer Systems Corporation
9055 Comprint Court
Gaithersburg, Maryland 20877

1-800-638-4041
(1-301-258-5060, in Maryland)

IBM is a registered trademark of International Business Machines Corporation. Referenced model numbers are the products of International Business Machines Corporation.

* Subject to PSC terms and conditions.

CIRCLE #50 on Reader Service Card

Securing Data On The VM System: A Rational Approach

By Steve Halperin and Steven Rev

The wonderful thing about IBM's VM operating system is that it creates an illusion that each user of the system has the computer all to himself. The vulnerability of his or her data is not an illusion, however, and the VM professional must face the fact that VM presents considerable opportunities for sensitive data to be exposed, altered and/or destroyed.

Trespassing, the act of unauthorized access to the computer, was once the "hot" security issue for mainframe computer users. Aside from the emotional issue of intruders using expensive system resources, however, there was little real damage to the owner of a computer if a hacker or DP professional cleverly bypassed security to spend a few minutes or hours on it. It was unnerving, but not the end of the world. Nevertheless, considerable efforts have been made to prevent unauthorized access to and use of mainframe computers by IBM in the systems software and RACF product and by independent vendors with products such as the widely used ACF2.

While trespassing remains an issue today, uglier motives for trespassing have significantly deepened the consequences of computer crime in the business community. Both outsiders and insiders (mostly insiders) are now stealing and altering computer data for personal gain or revenge. Computer crime statistics are becoming as depressingly familiar to everyone in data processing as are crime statistics in general to urban America.

Now DP professionals need to be concerned not so much with access as with the threat to the data that is maintained on their system and on computer tapes.

The toll taken by these insidious acts of computer crime weigh heavily on the

companies that are victimized in revenue loss, productivity setbacks, competitive disadvantage and lawsuits. Certain key individuals at the company are also paying the price. The company's directors may be financially liable for the misuse of the company's data, DP managers face dismissal for failing to adequately provide a secure environment for the company's data and systems programmers can have their records tainted. The old adage of "once burned, twice shy" does not apply in security disaster post-mortems. Often with 20/20 hindsight, senior management looks askance at the concept that security is not important enough to justify all-out preventive measures. A more applicable saying in the case of compromised file data is "heads will roll."

Just as damaging, though not as spectacular, is the practice of "snooping" where users access other individuals' files and view sensitive mail, memos, budget reports and so on that may contain proprietary information. The consequences of such actions range from ordinary embarrassment to serious leaks of information that can hurt internal operations.

The Need for Better Defense

What preventive measures can DP centers take, especially if they are running VM systems? Already familiar with products such as RACF and ACF2 from their MVS days, many VM professionals have implemented such packages on their VM systems in the hope or belief that the extra layer of access control will in fact protect the system's data.

And yet because of inherent weaknesses in VM itself, it is impossible to guarantee the protection of file data

through access control software. Even the makers of ACF2 admit as much in their general information manual stating that the lack of integrity in the CMS environment limits any security system from providing absolute file-level protection.

Access control is thus only the first step in providing a secure environment, but any installation that has gone to the expense and effort to implement such a system has surely done so because it has determined the need for totally securing its information. Given the limitations of VM, however, it becomes necessary to augment access control procedures with the only proven method of guaranteeing the security of file data on the VM system: data encryption as a stand-alone procedure. (See Figure 1.)

Data encryption renders data completely unintelligible to anyone who is not authorized to use it, even if they manage to gain access to it. Yet until recently, full-featured, stand-alone encryption packages have not been a viable solution for VM sites because of overhead considerations. VM1, a Los Angeles-based software vendor specializing in VM systems software, has developed a new product called SafeGuard to provide VM sites with an easy-to-use and inexpensive means of securing their file data, regardless of whether or not they are already using an access control package. For those who have considered the expense of access control justifiable for security reasons, the incremental cost of obtaining a viable stand-alone encryption product need no longer be a factor. Even IBM, realizing the limitations of access security, encourages the use of encryption.

The need for such products as an alternative or supplement to access control systems, becomes evident when consid-

A Boost In
VM Performance
So Amazing

*You'll think it's
Magic!*

It Is-CACHE MAGIC™

- ★ Increase system capacity without adding hardware
Access user files at processor speeds
- ★ Improve service to users
Eliminate batch overruns, speed online response times
- ★ See immediate results
Installs easily and works automatically

Call or write today for a free brochure.

SDI

1 (800) 537-1969 (415) 572-1200 In Europe: + 31 10 436 12 77
1700 S. El Camino Real, San Mateo, CA 94402

CIRCLE #76 on Reader Service Card

ering the fact that computer crimes can be so damaging. For example:

- Security Pacific National Bank of Los Angeles lost \$7,000,000 to someone who tampered with the bank's electronic funds transfer system.
- The Immigration and Naturalization Service was victimized by its own computer security official who used the system to create fake immigration documents which he sold to illegal aliens.
- An employee of the California Department of Motor Vehicles was convicted of contracting with an automobile thief to change the data in the department's vehicle ownership files to indicate that the thief was the real owner of the car he intended to steal.
- Volkswagen AG of West Germany suffered the manipulation of accounting records to cover up nearly \$259,000,000 in trading losses on foreign currencies.
- The J. Walter Thompson company discovered that the billings in one of its departments had been inflated by \$23,000,000.

The potential for more disastrous crimes is even greater when the types of information stored on corporate computers is considered. For example:

- Litigation in the case of IBM, Hitachi and Mitsubishi includes the charge that books containing IBM's business plans were stolen from that company. Many companies keep their business plans on their computers.
- Reader's Digest in Denmark and American Express in the United States suffered thefts of mailing lists that can be vital to the success of companies engaged in direct mail advertising. Mailing lists are commonly stored on computers and computer tapes.
- Other information stored on computers that comprise a company's major assets are computer programs, plans for new product introduction, manufacturing details for new products, client files, plans for mergers and acquisitions and other trade secrets.

If the unauthorized, benign use of a VM system is the only security concern of a company's senior management, then the company's data center can probably make do with VM's minimal built-in protection. However, if the company has already determined that expensive access control is warranted, then it also has data

that it needs to protect. To fully protect its data, both on the system and on tape, encryption is also required. To understand why, it is first necessary to understand the fundamentals of VM.

VM: The Return to Multi-User, Interactive Systems

The mainframe computer started out as single-user, dedicated device much like today's personal computer. In those days, however, computer resources were too expensive to be commanded by single users. So batch multiprogramming systems were developed at the expense of the single-user, interactive environment.

The return to interactivity began in the early 1960s, when a group at MIT developed the CTSS time-sharing system. While CTSS ran a conventional batch system that kept the computer busy at all times, it gave fast responses to interactive users doing programming and debugging.

Influenced by the success of CTSS at MIT, a group at the IBM Cambridge Scientific Center developed CP/CMS that came to be used on IBM's 360 series computers in the late 1960s. This evolved into VM/370 that became available for the 370 series in 1972 and VM/CMS in use today.

VM simulates the concept of an entire machine dedicated to one user bringing back the benefits of the single-user, interactive environment. Its two basic components are CP (Control Program) that provides each user with a simulated (or virtual) dedicated machine and CMS (Conversational Monitor System) that runs on a CP-created virtual machine as a one-user, interactive system. With VM, the system can be used at full capacity as with a batch system, but the economies of scale are achieved by making the system available to multiple interactive users.

The Other Side of the Coin

The advantages of VM in terms of multi-user, simulated dedicated access become disadvantages in terms of securing the system and its data against abuse. Everything about the concept of VM encourages ease of entry, ease of use and ease of access to computer resources.

There are two system attributes provided by VM that give the user considerable powers:

1. A DASD volume on which data files reside. The volume can be dedicated to one particular virtual machine or, usually, shared by many.
2. A virtual console that simulates a complete physical console to the real ma-

chine. Thus, the user can issue from the terminal commands such as:

- QUERY — to obtain information about the system and its resources
- LINK — to obtain the use of a DASD volume or portion of same
- SPOOL — to control the transfer of data files among users.

Because of the shared use of data storage devices and the user's ability to query the system, link to storage devices other than his own and control the spooling of data files, VM has built-in security problems. These problems have not gone unaddressed by IBM, however, as VM attempts to respond to each potential security issue:

- Access to the system is controlled by user IDs and passwords. A security journaling option (specified when the system is generated) records unsuccessful logon attempts and sends messages to the system administrator documenting the date, time, terminal, user ID and passwords involved.
- Privilege classes are assigned to each user that define the limits of system facility access. Certain commands and functions are associated with specific privilege classes and an attempt to execute a command that is out of bounds for that user's class will result in denial of command execution.
- Minidisks (the divisions by CP of a DASD volume) on which data files reside can be access controlled by password and can be made read-only, so that data cannot be altered on them when they are accessed.

Each of these precautions is accounted for in the VM Directory, a file containing descriptions of all of the virtual machines. There is a directory entry for each user, with ID and password, privilege class, dedicated minidisk assignment and shared minidisk assignment.

The controls placed on the abuse of the system are therefore related to access and the knowledgeable user who circumvents these controls by changing his or another user's directory listing (in the most basic of terms). This will gain access to areas of the computer to which he is not authorized to go and to command functions that he is not authorized to use.

CP operates in the real machine's supervisor state and the virtual machines run in the real machine's problem state. Each virtual machine can operate in either virtual problem state or virtual supervisor

**Would you spend
a few thousand dollars
to save millions?**

They did.

How much is a corporation's data worth?
Some say millions. Some say it's priceless.
Some don't know until they lose it.

The DATA RECOVERY SYSTEM (DRS) from Integrity Solutions is the world's leading VSAM data integrity system for IBM mainframe environments. With it, both CICS and batch VSAM data can be quickly and easily recovered in the event of inadvertent data corruption or loss. Redundant batch backups can be eliminated. And CICS journals and logs are automatically archived with the JOURNAL MANAGEMENT SYSTEM (JMS). Over four hundred of the world's leading IBM shops can't be wrong.

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



Integrity Solutions, Inc.

7921 SouthPark Plaza, Suite 200
Littleton, CO 80120
1-800-289-9900 or (303) 794-5505

CIRCLE #43 on Reader Service Card

©1988

state. When it issues a privileged command, an interrupt is generated. CP handles the interrupt determining the course of action depending on the virtual register and virtual state word maintained for that virtual machine. If the virtual machine is in virtual problem state, the interrupt is passed back to the virtual machine (denied). If the virtual machine is in virtual supervisor state, CP will simulate the execution of the privileged command, but only if the virtual machine is authorized to execute that command.

Penetration studies have found that one of the major security vulnerabilities of VM results from the complexity by which CP simulates the input/output interfaces. Penetrations are possible because of errors in simulating certain channel programs that can be written by a user and from a failure to anticipate certain effects made possible by the simultaneous operation of the CPU and input/output in the support of a user.

Furthermore, since the CMS environment permits the general user program to execute in supervisor state, to enter storage protect key zero, to execute privileged commands, to issue input/output commands and process interrupts independently of the CMS nucleus, the sophisticated user can potentially modify the CMS nucleus compromising CMS and OS file-level security.

CMS file security is basically divided into two components—one controlled by CP and one by CMS. The component controlled by CP that is at the minidisk level, provides either read, write or multiple read/write access to any particular minidisk. The CMS component attempts to add an additional level of security by preventing a user with read/only access to a minidisk from seeing files designated as filemode zero. The failure of this method of protecting files is that the CP control is not aware of any of the CMS level file protection schemes. Therefore, when CMS is not operating or is bypassed or is modified as previously mentioned, restricted files can be obtained for misuse.

This is why VM's built-in security measures are insufficient for protecting file data and why any access control system is limited in providing absolute file-level protection.

The answer to this deficiency in VM has traditionally been to add more layers of access control that only makes a user's efforts to obtain the use of data more difficult rather than impossible. Unfortunately, the level of difficulty serves to lull

many data center personnel into a false sense of security. The crime statistics belie the best intentions of the package's designers, particularly when companies are victimized by their own employees. And no access system protects data stored on tape.

The argument has also been set forth by the administrators of access control packages that audit trails serve as a deterrent to computer crime. Since a record is kept of all activity on the system, anyone committing an unauthorized act would be found out. The truth is that audit trails deter computer crime as much as the death penalty deters murderers. And even if detection of the crime was an issue for the potential criminal, according to the Na-

ideal procedure for the business community at large, because it solves a serious problem with shared systems such as VM. Even without the threat of outside intruders gaining access to data files, every system has at least one person designated as having unrestricted access to all system files. Users (and companies) may be reluctant (wisely) to store really sensitive information. Encryption enables users (and companies) to trust their most valuable information to mainframe computers and thus have more respect for the mainframe and the data center.

Furthermore, if encrypted tapes or disk packs are stolen from the facility or enroute via mail or courier service, they are useless to the thieves. Encryption is the

FIGURE 1

How Secure Is Your VM System?

SECURITY ISSUE	VM SYSTEM	ACCESS SECURITY	ENCRYPTION
Protects data at file level	No	No	Yes
Protects data on tape	No	No	Yes
Control of resources (computers, disks and so on)	Some	Yes	No
Allows user-level security control	No	No	Yes
Avoids dependency on passwords	No	No	Yes
Easy to use/implement	Somewhat	No	Yes
Ensures data is secure in event of system crash	No	No	Yes
Provides audit trails	Some	Yes	No
Insulated from central "security czar" who can manipulate system	No	No	Yes
Invulnerable to inside user who can bypass security procedures	No	No	Yes

tional Center for Computer Crime Data there is little reason to expect that anything will be done about it. The odds are one in twenty-two thousand that a computer criminal will go to jail, because the definition and prosecution of computer crimes are so nebulous and ineffective in the court system. Regardless of the consequences of the crime to the criminal, however, the consequences to the victim are certain. A better solution is in order.

Encryption: The Key to Safe Data

Data encryption was once considered the sole interest of the military and diplomatic communities. The National Bureau of Standards adopted a Data Encryption Standard (DES) to apply to all encryption of sensitive government information. Today, however, encryption is being considered more and more as an

only way to secure data once it is offloated from the system for backups or transfer.

Encryption can also be used to limit the use of functions on the system, not merely to secure data files. If program modules are encrypted, they are rendered useless and only those people authorized to use them would be given the encryption key (the encryption/decryption of modules could be automated in an exec).

In a symmetrical encryption scheme, such as DES, the same key is used for both encryption and decryption. The key is a string of characters that is used in conjunction with an algorithm, a fixed set of rules in a mathematical formula. The intruder must know the key in order to decrypt the file, even if he knows the algorithm. For further protection, more than one key can be used to encrypt files.

See Data page 75

SYSTEM DEVELOPMENT UNDER CICS



150 ISPF
Users/4381*

SYSD[®]

- TSO/ISPF like Edit and Browse
- Screen Management
- On Line Report Viewing
- Report Printing
- DASD Management
- CICS Display and Debug Utilities
- Full Security
- VS1, MVS, MVS/XA with CICS

Over 450 customers worldwide.

*150 ISPF/SDSF users on a 16 MB4381.
INTERESTED BUT SKEPTICAL? Ask for the story of a major company who's solution was recently featured in the 4300 Journal.

SYSD[®] \$27,000

PANVALET[®] Option \$5,000

PANVALET[®] is a registered trademark of Pansophic Systems, Inc.



H&W COMPUTER SYSTEMS, INC.
P.O. Box 15190 • Boise, Idaho 83715-0190
(208) 385-0336 • 1-800-338-6692
FAX: (208) 342-5219

CIRCLE #154 on Reader Service Card

Manager from page 20

have resolved the problem. If the user is not sure how the problem occurred, you can use the scroll feature to review the user's previous screens and determine what actually took place.

Messages

An important benefit of the message feature is the ability of the systems programmer to inform users of changes in the operation of the network. This is done with the broadcast message. Look at the

broadcast message feature from an operations perspective.

The ACF/VTAM systems group has informed operations that they want to reload the 3725 at noon to correct a network performance problem. There are production lines on the 3725 that must be notified of the outage.

First, using the broadcast feature of the session manager, the operator can notify all active users of the 3725 that they will be shut down at 11:45 for 15 minutes to reload the NCP. Second, the operator can

set up a log-on message for the users who log on after the broadcast message is sent. When those users log on to the session manager, this message notifies them that the 3725 will be down. The third step is to schedule a timed message for delivery at 11:30 to the 3725 users reminding them of the 11:45 NCP change. At 11:45 the operator can feel confident that most of the users have been notified of the pending shutdown and he can proceed to reload the 3725.

Usually, the problem of informing users who access several subsystems of a pending change comes in locating them. This has been solved through the use of the session manager's broadcast services. A session manager enables you to set up log-on, broadcast and timed messages targeted at terminals, groups of terminals, users, groups of users or users of specific subsystems like TSO or CICS. This flexibility allows systems and operations personnel to communicate changes easily and quickly which minimizes the effect on the users.

All of these features provide operations and systems personnel with the tools they need to help end users when problems occur. These tools can be used to teach classes, conduct application reviews and run on-line meetings. Help desk features are tools that help the MIS department contribute to the company's goal of creating a more productive and profitable business environment.

Network Performance Features

The last category of features provided by software session managers is made up of the network performance items. Among these features are the statistics related to network traffic and user response time. This category also includes the 3270 data stream compression techniques that help in the reduction of needless character transmission in the network.

Utilization Statistics

Session managers are an ideal place to collect statistics on end-user response times. The session manager sits in the middle between the user's terminal and the subsystem. Therefore, the session manager can measure the time it takes for a transmission to complete to a terminal or the turn-around time associated with different subsystems. The session manager can also keep counts on the amount of traffic that a particular user is generating and the overall connect time of that user to different subsystems. Figure 3

See Manager page 50

The VSAM Space Manager™


- Pools DASD for VSAM
- Eliminates production ABENDS due to lack of space
- Automatically controls allocations based on your shop standards
- Keeps test files from cluttering production packs
- Reports on VSAM DASD use and trends
- Makes IDCAMS JCL device and volume independent

For more information call

800-638-9254

MVS only.

People



SOFTWARES, INC.

Know Surprises

Now when you make changes in your CICS systems, you don't have to wait for users to determine that they work. Or even worse, that they don't.

With Compuware CICS PLAYBACK, you can create completely accurate scripts for testing.

Because CICS PLAYBACK enables you to capture real-life transactions that can be played back to test. Designed for IBM mainframe environments, it has the architecture to execute all five levels of testing: High-volume, Regression, Integration, Concurrency,

and Interactive Unit level.

Avoid the high cost of downtime in your CICS network. Whether you're measuring the effects of a new applications package, or configuration changes, know the results without having to pay for the results — with CICS

PLAYBACK.

For more information on CICS PLAYBACK, CICS Abend-AID and CICS dBUG-AID, write, or call us at: 31440 Northwestern Hwy., Farmington Hills, MI 48018-5550, 1-800-521-9353. In Michigan, (313) 737-7300.



COMPUWARE
Because experience is everything.



Candle

**The Power of
Automated Operations**



Is Here.

Power You Trust Behind the Master Console

It takes power to automate. Not just iron and a little software. But a force behind the automation tool that can help when things get tough. Candle, the company that defined MVS performance with OMEGAMON[®], has entered the automated operations field with AF/OPERATOR[™].

Energize Just What You Need

Candle's family of automation products is flexible enough to meet your individual needs—from simple message management to remote operations. It's up and running in an hour, handling basic message traffic and doing flawless IPLs at warp speed.

Or Blast Off Into Full Automation

When you're ready to automate recovery or complex procedures, AF/OPERATOR enhances your growth. With an automated operations product from Candle, you're uniquely positioned to handle whatever you need in the engine room, including automated availability and performance when the time comes.

Streamline the Engine Room

Basic message management is the first step. And with AF/OPERATOR's sample library, that first step is a snap. Yet, your horizons are never limited because you have the power to tap into fully integrated automated operations with the ability to control MVS and its subsystems—IMS, DB2, CICS, and JES.

With High Performance

System overhead isn't an issue when you can have less than 1% CPU consumption. Nor is scheduling when you can make changes to the system on the fly, rather than waiting for an overnight 'gen'.

And Support From the Mother Ship

When you take a step into the dimension of automated operations, you want the best ground crew behind you. And Candle's support has received DataPro's highest rating. This support extends beyond round-the-clock customer service and education. Candle spends \$30 million a year in R&D so you won't ever be left behind.

So whether you're going to the outer limits of automation or just streamlining the engine room, Candle is ready to join your crew. Let automated operations beam down into your data center. Call Terry Forbes today at **(800) 843-3970**.

!Candle[®]

Candle Corporation
1999 Bundy Drive
Los Angeles, California 90025

Copyright ©1988 Candle Corporation. All rights reserved.

CIRCLE #120 on Reader Service Card

Manager *from page 46*

shows the relationship of terminals and applications when an ACF/VTAM session manager is present.

Some session manager products include a feature that reports on system

availability. Because the session manager is logging on and off of several applications, it is a logical place to measure the number of successful log-on attempts providing a measurement of availability.

3270 Data Stream Compression


Session managers are an ideal place to do 3270 data stream compression. The compression techniques can be applied without changing the subsystems or user

See Manager page 73

Multiple-Session Manager Products

Product Name	Vendor	Brief Product Description	Operating Environment	For More Information
Multiterm	BlueLine Software, Inc. 1500 S. Lilac Dr., Suite 340 Minneapolis, MN 55416 (800) 826-0313	The Multiterm family of products provides comprehensive session management and HELP DESK facilities for VM, MVS and DOS/VSE users. HELP DESK facilities can be used by authorized technical personnel to identify and resolve problems experienced by terminal users.	MVS, VSE, VM/SP, VTAM	CIRCLE 300
CL/Supersession	Candle Corporation 1999 Bundy Dr. Los Angeles, CA 90025 (213) 207-1400	CL/Supersession is a multi-session manager that supports an unlimited number of concurrent users and sessions while maintaining maximum efficiency and throughput. CL/Supersession is built around a Dialog Manager that provides for complete customization.	MVS, MVS/XA, VTAM	CIRCLE 301
CA-VMAN	Computer Associates International, Inc. 711 Stewart Ave. Garden City, NY 11530 (800) 645-3003	CA-VMAN is a network and session manager that allows single or multiple sessions by terminal or user. It provides a single entry point to VTAM, extensive security interfaces, menu selection of network applications and generates VTAM and NCP parameters.	MVS, VTAM	CIRCLE 303
TPX (Terminal Productivity Executive)	Duquesne Systems, Inc. Two Allegheny Center Pittsburgh, PA 15212 (412) 323-2600	TPX improves end-user productivity by allowing simultaneous connection to several on-line applications and minimizes consumption of network resources. Users can switch between multiple applications with a single command or PF key.	MVS, VM, VTAM	CIRCLE 304
TUBES	Macro 4 Systems Software Brookside Plaza Mt. Freedom, NJ 07970 (800) 223-0414	The TUBES family is a group of products that provides all users logged on to an easy means of accessing all of the facilities that may be required by a user in a VM environment or all applications required in the VTAM environment. Each TUBES user may have concurrent access to up to 12 systems while TUBES VTAM users may have concurrent access to up to 24 different applications.	MVS, VSE, VM, VTAM	CIRCLE 305
Televue	Pansophic Systems, Inc. 709 Enterprise Dr. Oak Brook, IL 60521 (800) 323-7335	Televue provides advanced SNA session management capabilities for concurrent terminal sessions, global network security and session pooling in VTAM networks. Televue resides as a VTAM application in one or more MVS post processors.	MVS, MVS/XA, VTAM	CIRCLE 306
SIM/Session II	Simware, Inc. 20 Colonnade Rd. Ottawa, Ontario, Canada (613) 727-1779	SIM/Session II is a session manager that supports multiple VTAM applications concurrently from a single physical terminal. It enhances user productivity by providing multi-session support, reduces hardware cost and helps achieve a higher return on existing resources.	MVS/XA, VTAM	CIRCLE 307
NET-PASS	Software AG of North America 11800 Sunrise Valley Dr. Reston, VA 22091 (800) 843-9534	NET-PASS is a VTAM session manager that provides easy menu access for single session users, low overhead session switching for multiple session users and input and output optimization techniques for remote users.	MVS, MVS/XA, VSE, VM, VTAM	CIRCLE 308
PIE	Technologic Software Concepts, Inc. 16842 Von Karman Irvine, CA 92714 (714) 863-7080	PIE's "preferred environment" technology for TSO, CICS and VM provide access to multiple VTAM applications. PIE's MultiTSO and MultiCICS components allow a user to have multiple, parallel applications within TSO and within CICS without any VTAM overhead.	MVS, VSE, VM, VTAM	CIRCLE 309
MULTSESS	Westinghouse Management Systems Software P.O. Box 2728 Pittsburgh, PA 15230 (800) 348-3523	MULTSESS, designed to be transaction efficient, provides users with concurrent access to multiple VTAM applications. MULTSESS features fully customized dynamic selection menus, teacher/pupil mode, data transfers between sessions and integrated on-line administration panels.	MVS, MVS/XA, VM/GCS, VTAM	CIRCLE 310

MOST DASD PERFORMANCE SOFTWARE TELLS YOU WHAT'S GOING WRONG.



WE TELL YOU HOW TO MAKE IT RIGHT.

Until now, DASD tuning was a painstaking task-requiring hours of statistical analysis just to figure out what was wrong. Correcting problems often involved guesswork, trial and error and just plain luck. DASD ADVISOR from Boole & Babbage tells you exactly what's going wrong and how to make it right.

DASD ADVISOR is an EXPERT system-based DASD tuning tool that eliminates the need to wade through piles of performance

statistics. It analyzes the performance of your entire DASD subsystem, from individual data sets, through hundreds of devices, controllers and channels. It identifies data bottlenecks, then makes specific tuning recommendations. All in concise English. So you have what you need to improve DASD performance. And time to solve other system performance problems.

For a free demo diskette that shows you how DASD ADVISOR can help you make it right, call Marty Johnson. In California: 800-624-5566. Outside California: 800-822-6653.

Boole & Babbage, Inc., 510 Oakmead Parkway, Sunnyvale, California 94086.

Boole & Babbage 

The Performance People

CIRCLE #160 on Reader Service Card

A PROLOG-Based Expert System For Tuning MVS/XA

By Dr. Bernard Domanski

In his acceptance speech for the A.A. Michelson award in 1985, Dr. K. Mani Chandy said that it takes some ten to fifteen years for research in the universities to become practical technology for industry. The field of Artificial Intelligence (AI) is a prime example. The last fifteen years have seen AI flourish in coming from the university sector to industry. AI languages such as LISP and PROLOG have been touted as the languages of "choice" for fifth-generation computing. The Japanese have chosen PROLOG as the machine language of their logic processor for their fifth-generation system. Expert System shells such as OPS5 and M.I can help simplify and speed the development of these systems.

Considering Computer Performance Evaluation (CPE), some prior exploratory work has been done. Vendors are now emerging with knowledge-based systems for CPE (ISS, Boole & Babbage).

Although AI has had a long history, the application of AI to CPE is still immature. In CPE, many performance analysts have developed a set of "rules-of-thumb" for tuning large systems. Typically, this set takes five years to develop. As there is a tendency to promote quality personnel within a corporation, it is critical to extract the set of rules from the performance "expert" before (s)he is promoted to a new position. In essence, a "technology transfer" must occur.

This article will discuss some of the issues involved in building an expert system that embodies tuning rules for IBM's MVS/XA operating system. To understand the components of an expert system and their functions, PROLOG on an IBM PC (Turbo-PROLOG from Borland Inter-

national) was chosen as the development environment. The article will begin by defining the key concepts about expert systems, knowledge engineering and knowledge acquisition. The reader will be given a brief overview of PROLOG which will explain how an inference mechanism was developed. Finally, the article will describe the expert system that was developed and additionally will provide a set of key issues that should be addressed in the future. The overall objective is to provide new insight into the application of AI to CPE.

What Is an Expert System?

An expert system is a knowledge-based reasoning system that captures and replicates the problem-solving ability of human experts. Expertise consists of knowledge about a particular area or domain, understanding of domain problems and skill at solving some of these problems.

Domain knowledge, rather than the complexity of formal reasoning methods, is the key to solving difficult problems. Knowledge can take many forms:

- Specific inferences that follow from specific observations
- Necessary and sufficient conditions for achieving a goal
- Probable causes of symptoms.

Knowledge can be represented in sev-


eral ways such as rules, frames and logical predicates. In most current expert systems, knowledge is expressed in rules and these rule-based systems will be discussed here. The skill of these systems increases at a rate proportional to the enlargement of their knowledge bases. They usually explain their conclusion by retracing their lines of reasoning and translating the logic of each rule used into natural language.


There are a variety of reasons for building systems that capture human expertise:

- Experts retire or are promoted taking their knowledge with them
- An expert's time is not particularly well spent answering user's questions
- Expertise may be scarce
- Expertise may be expensive to deliver
- Experts are not always consistent.

Figure 1 shows a diagram of the most basic components of an expert system. The knowledge base is usually a collection of rules about the problem domain supplied by the expert. Problem solving knowledge is usually used in consultations between the expert system and the end user or with other programs. The working memory keeps track of what the system knows during each consultation. The

FIGURE 1





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

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

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

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

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



Levi, Ray & Shoup, Inc.

Specializing in Computer Systems Software

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

Telex (510) 60-00675



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

inference engine looks in the working memory and the knowledge base to see what is true at any given time to resolve queries.

Consider a simple knowledge base with only two rules:

1. Rule 1: IF *a* THEN *b*
2. Rule 2: IF *b* THEN *c*.

Each rule has a premise clause (IF) and a conclusion clause (THEN). If a user poses the question "IS *c* TRUE," the inference engine looks through all the rules in the knowledge base and tries to find one that would conclude about *c*. It finds rule two: IF *b* is true, THEN *c* is true. Now, if the inference engine could discover whether *b* was true, then it would know whether *c* was true. So a query is made of working memory, "Is *b* true?" Again, the inference engine examines the rules in the knowledge base, this time looking for a rule that would come to a conclusion about *b*. It finds rule one and this leads to asking the question, "Is *a* true?" If the user responded yes, then the inference engine would:

- erase the query about *a*
- place the fact that *a* is true in the working memory
- use Rule 1
- place the fact that *b* is true in the working memory
- use Rule 2
- place the fact that *c* is true in the working memory
- resolve the original query and
- tell the user "Yes, *c* is true."

This is a simple example of a backward chaining system. The inference engine works backwards through the conclusions of the rules, trying to determine whether the premises are true. In a forward chaining system, facts would be placed in the working memory and the inference engine would report any conclusions it could to the user. For instance, if "*a* is true" were entered, the system could report that both *b* and *c* were true.

Forward chaining was used in our PC/PROLOG Expert System for MVS/XA. Backward chaining systems tend to be goal driven which means they work backward from some specific goal (*c* above). Forward chaining systems tend to be data driven which means the system is primed with certain facts and draws any conclusions it can.

Other information must reside in the knowledge base in addition to the rules. For example, consider the user interface. How does the expert system know what

question to ask the user when that becomes necessary? This information is stored in the part of the knowledge base that describes the *as* and the *bs* or the traits characteristic of a given problem domain. For example, the trait USERS may carry with it the prompt, "On average, how many users are connected to the system?"

However, when many rules are present in the knowledge base, complex consultations can occur. It can become difficult to predict what questions will be asked and in what order unless some precautions are taken. The builders of the famous MYCIN medical expert system found that doctors were somewhat bewildered when the consultation questions jumped from one topic to another: "What is the patient's age?" "What was the result of test X?" "What is the patient's name?" and so on. It becomes necessary to use more control information. For instance, the system could be instructed, "Ask about traits *a*, *b* and *c* first, then try to find the value of trait *d* by forward chaining, then ask about traits *e* and *f*, . . ." This consultation and control information is also kept in the knowledge base. As will be shown later, the PC/PROLOG expert system for MVS/XA employed this concept in its structure of its knowledge base.

It is sometimes necessary to keep even more information in the knowledge base such as special purpose algorithms or methods needed to retrieve information from databases that reside outside the expert system. Thus the task of knowledge engineering includes more than eliciting rules from the expert. All the traits (such as *a* and *b* above) and their characteristics and accompanying control information must be identified as well.

Advantages of a Knowledge-Based System

A natural question asked about expert systems is, "What are the advantages of using a knowledge-based system?" In general, knowledge-based systems extend the speed, accuracy, consistency, availability and affordability of computation from clerical to intellectual applications. Specifically, the advantages can be seen in the following areas:

- Behavior: users can ask "Why?" questions to show why certain decisions were made. This is in direct contrast to a traditional data processing system in which the only way to verify the system's reasoning is to simulate the system's behavior.

- Knowledge independence: a piece of knowledge (a rule) can be examined and easily changed somewhat independently from other pieces of knowledge in the system.
- Separation: knowledge is kept separate from the reasoning process that uses it. Building an expert system consists of changing knowledge rather than program code. The expert and the knowledge engineer can examine and change this knowledge since it is ideally represented in an English-like format. Because of this separation, rule-based systems can be incrementally developed with steady performance improvements.

Knowledge Engineering/ Knowledge Acquisition

What is knowledge engineering? This is the term given to the extraction, description and computerization of expert knowledge. Knowledge consists of descriptions, relationships and procedures in a domain of interest. The major task in building an expert system is the transfer of knowledge from its source into the expert system. This is called knowledge acquisition.

Potential sources of knowledge include human experts, textbooks, databases, examples, case studies and personal experience. Within CPE, nearly all these sources can be used to elicit performance knowledge. The prototype system was developed using PROLOG-used knowledge that was extracted from technical papers, IBM manuals and personal experience.

The process of building a comprehensive expert system can be described as an iterative, cyclic interaction between the expert, the knowledge engineer and the system itself:

1. The expert tells the knowledge engineer what rules to add or change
2. The knowledge engineer makes changes to the knowledge base
3. Several tests are run for consistency checking
4. If any problems arise, they are discussed with the expert and we go back to step 1
5. The expert runs the modified system on new cases until new problems are discovered
6. If no problems are encountered in a substantial number of cases, the expert stops checking the system.

The stages of expert system knowledge acquisition roughly correspond to systems

IMAGE COPY

Faster Image Copies

IMAGE COPY PLUS provides exact image copies of your DL/1 data bases *much faster* than the IMS utilities you are now using. IMAGE COPY PLUS *reads* the data base faster, *writes out* the image copy faster, and *processes multiple tasks asynchronously* to further reduce wall clock time.

Virtually No Data Base Downtime

The Incremental Image Copy feature allows you to produce an image copy without accessing the data base. And the Online Image Copy feature keeps your data bases available to users *without* affecting system response time.

Reduced Copy Tape Consumption

Three compression options reduce the amount of data to be written out—further reducing the time needed to create image copies and greatly reducing tape requirements.

PLUS

For more information or to begin a 30-Day-Plus Free Trial, complete and return this coupon or call BMC Software, Inc.:

1-800-841-2031

in the USA or 713-240-8800
1-800-231-2698 in Canada

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

BMC
SOFTWARE

MJ

Contact me about a 30-Day-Plus Free Trial of IMAGE COPY PLUS.

Contact me with more information on IMAGE COPY PLUS.

Name _____

Title _____

Company _____

Address _____

City _____ State/Prov. _____ Zip/PC. _____

Phone _____

analysis and software engineering according to Hayes-Roth (Rule-Based Systems). In this acquisition model, an expert system grows by proceeding from simple to increasingly more difficult tasks and by incrementally improving the representation and organization of the knowledge. We outline this acquisition model below.

- **Identification:** identify the problem area and define its scope while determining the resources needed and deciding on the goals and objectives of the system. Questions that must be answered include, "What class of problems will the expert system be expected to solve?" "What data is required?" and "What are important terms and their interrelationships?"
- **Structure:** key concepts, relations and information-flow characteristics needed to describe the problem-solving process are defined. Related questions might include, "What is given and what is inferred?" "How are objects related?" and "What is the information flow?"
- **Formalization:** concepts and relations are mapped into a formal structure suggested by the expert system. Related questions include, "What kinds of questions should be asked to obtain data?" or "What is the cost of data acquisition?"
- **Implementation:** a knowledge engineer combines and organizes the formalized knowledge to make it compatible with the information flow characteristics of the problem resulting in a prototype that can be executed and tested.
- **Testing:** the performance of the prototype is evaluated and revised to conform to the standard defined by experts in the problem domain. Typically, this is an iterative process: an expert evaluates the program's performance while the knowledge engineer revises the knowledge base.

It is beyond the scope of this article to examine knowledge acquisition techniques in detail. Note that most of these techniques require that knowledge be elicited from information sources and first placed into an information base. From there, it is analyzed and organized into knowledge bases and then used with an expert system. The knowledge must be tested for necessity and sufficiency and the knowledge bases must be refined accordingly. Techniques in knowledge acquisition often have roots in psychology.

Our PC/PROLOG-based expert system

for tuning MVS/XA faced the same problem of knowledge engineering/knowledge acquisition. Though we chose as our initial knowledge source written material (that is, papers and IBM manuals), we first addressed questions of identification. We felt that the most common types of tuning problems for MVS/XA are I/O related; thus, the class of problems for our first prototype was limited to I/O and paging/swapping. We next needed to identify the data that would be required and this was inferred from our information sources. For example, we decided that it would be important to know whether or not dedicated page volumes were being used. Overall, the objective of our system was to gather information about the presence/absence of performance "symptoms," analyze these to form a "diagnosis" and then suggest "treatment."

In defining the structure, we grouped certain traits together that are related to specific problems. For instance, values for pend time, connect time, disconnect time and IOS (IO Subsystem) queue time are related to poor I/O response time. We could derive a value for I/O response time from these traits. The procedures for finding a problem could then be informally described by grouping the key traits together, inferring any additional information and finally describing (informally) the relationships (rules) between the traits and the problems to be solved.

The formalization of structure into a more formal representation is usually made easy when an expert system shell is used. The format of rules is specified by the shell, so the transformation does not have to be difficult. In the case this author worked on, no such shell was used. The problem was compounded because a new shell had to be developed. In a later section of this article, the steps that were taken to build a shell and to determine a "language" for expressing CPE rules will be described.

Implementation implies that the system be built. This process will be described later. However, an interesting question about the focus of rules comes up. We might build a rule like, "If page delay time is greater than 100 ms, then you should examine the page/swap dataset configuration." Here, the presence of the symptom (page delay time > 100 ms) implies a potential diagnosis which would have a corresponding treatment (that is, consider dedicated page volumes). We will call this a forward leading predicate. Now consider the following: if a treatment were

"use dedicated page volumes," then the expert system could ask, "Are you using dedicated page volumes?" If the answer was yes, then suggesting this as a possible treatment is redundant and confusing. This is an example of a backward looking predicate. The presence of a potential treatment would imply that the expert system should look for other treatments or reply that the system does not know a treatment for the problem.

Testing our prototype consists simply of letting an "expert" use the system and incorporating any additional knowledge the expert gives into the system. For this, our prototype was given a "why" facility: the ability for the user to ask how a particular diagnosis was reached. In addition, the prototype was given an editing capability over the contents of the knowledge base. Using commonly used PC editing keys, the expert could add, delete or modify not only the rules in the system; but also, the control mechanism and the contents of the help, diagnosis and treatment messages that the system uses. In short, the expert is given full control of all of the knowledge that the system has during the entire testing process. After testing, these features were left in the expert system.

PROLOG

Over the past two decades, the price of hardware has dropped dramatically while the costs associated with software development have now become the dominant portion of a total system budget. This rapid rise has influenced the development of new programming tools that simplify system development. PROLOG is the result of years of effort in this area. PROLOG was developed at the University of Marseilles, France, by Alain Colmerauer in the early 1970s. It was designed to be a convenient tool for PROgramming in LOGic (hence the name). It can be more powerful and efficient than many traditional programming languages (that is, PASCAL, FORTRAN and so on.)

A PROLOG program uses a description of a problem's facts and rules and then finds all possible solutions to the problem. In PROLOG, a programmer describes what must be computed, rather than how the computation should be carried out. This declarative (rather than procedural) approach eliminates well-known errors that are common in other programming languages (that is, one too few iterations in a loop). So aside from some initial declarations, a PROLOG program

It's not a High-speed Printer Unless it's High-speed on the Tough Jobs!



Introducing the NBS 3840 Laser Printer for a TRUE 40 pages per minute!

Mix text, graphics and forms, in any combination, and the NBS Southern model 3840 will still give you 40 pages of pure output per minute!

NBS Southern's model 3840 gives you forty honest pages per minute — form, text and graphics — all on a single pass on standard cut sheet paper at 300 dots per inch.

If you're tired of watching output speed degrade as forms and



graphics get more complicated, it's time to look at NBS Southern's 3840. It houses a true 57 resident fonts, with optional expansion to over 150. In fact, it supports the entire Bitstream® font library, and fonts may also be loaded directly from the host computer.

With the addition of NBS Southern's PAGEWARE™ Forms Description Language (FDL), you can easily merge text with complex documents, logos, signatures and other special graphics.

For computer printing requirements up to 200,000 pages per month, the NBS Southern 3840 is the ideal solution. It improves accessibility to the mainframe

from local terminals, it offers customized printer performance to meet local user needs, it eliminates the cost of pre-printed forms, and it solves the delivery problem inherent in moving printed materials from the printing site to a remote utilization point.

For detailed specifications and pricing information, contact NBS Southern, Inc., 11451 South Belcher Road, Largo, Florida 34643. Telephone 813/541-2200; outside Florida 800/327-5602; FAX 813/546-8042.

NBS
Southern, Inc.

consists of a set of facts (that is, john likes mary, tom likes sam) and a set of rules (that is, jean likes X if tom likes X). PROLOG deduces that jean likes sam. Goals like "find every person who likes sam" can be asked of a PROLOG program. PROLOG uses a built-in backtracking mechanism that, once one solution has been found, causes PROLOG to reevaluate any assumptions that were made to see if some new variable values will provide additional solutions.

To better understand the syntax and operation of PROLOG, consider the following sample program:

```
clauses
watches(bill,bob).
watches(john,jane).
watches(fred,felicia).
watches(mitch,bill).
watches(brenda,greg).
watches(bob,bob).
watches(fred,greg).
watches(bill,phillipe).
```

The clause "watches(bill,bob)" is the fact "bill watches bob." Thus, if we pose as a goal "watches(mitch,bill)," PROLOG would reply "true." "Watches(bill,jane)" would result in a reply of "false." Variables begin with upper case characters; thus, a goal of "watches(john,Who)" (analogous to asking Who does john watch?) causes a reply of "jane." But the goal of "watches(fred,Who)" results in both "greg" and "felicia." The "_" character is called the blank variable: a goal of "watches(fred,_) is analogous to asking "does fred watch anyone?" This would result in the reply, true.

If we add the rule "happy(eric) if watches(brenda,eric)" (that is, eric is happy if brenda watches eric), the goal "happy(eric)" ("is eric happy") would be satisfied. The rule "nervous(Who) if watches(bill,Who)" would be satisfied with "bob" and "phillipe." This is analogous to asking, "Who is nervous?" and answering using the rule "someone is nervous if bill is watching them." Thus, PROLOG looks through the already established facts and concludes that since bill watches both bob and phillipe, both bob and phillipe are nervous.

Though there are many interesting features of the language (the reader is encouraged to examine the cut operator, as well as backtracking), the most commonly used data structure is known as a list. Analogous to arrays, lists are collections of objects (elements) of the same type, separated by commas and placed inside square brackets. For example, [1, 2, 3] is a list of the integers 1, 2 and 3. PROLOG can manipulate a list by dividing it into two parts: a head and a tail. The head is the first object on the left in the list. The tail is the rest of the list. Syntactically, a PROLOG list is written with a vertical bar separating the head from the tail [head | tail]. Lists are treated just like other objects in PROLOG and PROLOG has built-in facilities (predicates) for manipulating the elements of a list.

Development of the Expert System

The reader should, as this author did, feel comfortable with the concept of an expert system, as well as some PROLOG basics. The discussion now turns to building the system itself. Consider that this author is (by definition) a programmer of procedural languages (FORTRAN, PL/I, C and so on.) Thus, PROLOG is some-

YOU SHOULD BUY THE BEST SCHEDULER, REGARDLESS OF COST.

In today's complex DP environment, you need the best automated job scheduler you can find.

You need Zeke, the Scheduler That Works. In a recent Datapro survey, users rated Zeke higher than any competitive job scheduler.

Unlike the others, Zeke is designed to solve the problems of today's data centers, not yesterday's batch environments. Only Zeke can open and close online files. Automatically reply to job messages. And bring the online system up and down. And that's important. Because job scheduling involves much more than simply submitting jobs.

Zeke is also the easiest job scheduler to learn. The easiest to implement. And it works with both MVS and VSE.

For a free copy of Datapro's independent analysis of job scheduling software, call Altai Software today. You'll find out why Zeke is the best job scheduler in the world—at any cost.

ZEKE
The Scheduler that Works.

what unnatural and causes some conceptual problems that have to be overcome. For example, consider the following rule: bad I/O response time exists if pend time is greater than or equal to 10 percent of the sum of disconnect plus connect time and IOS queue time is greater than the sum of pend time, disconnect time and connect time. In PROLOG,

```
bad_resp( IOSQ, PEND, DISC, CONN) if
  PEND >= (DISC + CONN) * 0.10 and
  IOSQ > (PEND + DISC + CONN).
```

This simple rule implies that our PROLOG program should:

- Prompt the user for the value of each variable (IOSQ, PEND and so on.)
- Provide help with each prompt should the user not understand what is being asked for
- Validate each resulting reply — here, with a simple numeric test
- Pass the replies as parameters to this rule that does the evaluation (bad_resp(IOSQ, PEND, DISC, CONN))
- If the rule is satisfied, provide the user with the diagnosis (poor I/O response time because of these symptoms)
- Look up a treatment for this diagnosis.

Our first version of the system has the rule (knowledge) coded in PROLOG; thus, it is part of the reasoning process. Given all the rules to be coded, plus the help, diagnosis and treatment messages, the resulting program was rather long and strongly resembled a traditional procedural program.

Recall that an objective of an expert system is to separate the knowledge from the reasoning process. Our first version mixed the knowledge and reasoning process together. Thus, we examined the rules a little closer, since they were the driving force behind the design of the program. When invoked, rules result in true or false. Thus, if the prompts were structured as true or false questions, the inference mechanism could be greatly simplified. For example, consider the following:

```
demand_paging_problem(P1, P2) if
  ask(i1, P1) and
  ask(i2, P2).

ask(Index, Prompt) if
  prompt(Index, Prompt) and
  write("Is", Prompt, "?") /*an ordinary write */
  and readln(Reply) /*read the user reply */
  and yes(Reply). /*validate the reply */

prompt(i1, "Demand Paging High").
prompt(i2, "storage isolation used for IMS
or CICS").
```

The inference mechanism "ask" is now

divorced from the knowledge of what prompts should be used. The prompts themselves are associated with specific traits (i1, i2). For illustration, we have included the prompt clauses within the program, but the collection of prompt clauses is ordinarily kept within a separate, editable file.

Note though that the control information (ask about trait i1 then trait i2 relative to demand paging) is still part of the PROLOG program. Thus we need to carry

this "separation" of knowledge from reasoning one step further. We create in a separate, editable file a knowledge base of the following form:

- prompt(trait, prompt messages).
- conditions(symptom, [list of traits]).
- depends(symptom, [list of traits], trait).
- diagnosis(symptom, [list of traits], treatment).
- treatmsg(treatment, treatment messages).

EVEN IF IT COSTS YOU LESS.

Just because a job scheduler costs more doesn't mean it does more. In fact, the best scheduler on the market—Zeke, the Scheduler that Works—will actually cost you less.

Not only when you buy the software. But also when you implement it. Because Zeke doesn't require you to make massive changes in your operations. Instead, Zeke adapts to you. So installation takes hours instead of days. And implementation takes weeks instead of months.

Once Zeke is fully in place, it can actually pay for itself in a matter of months. By cutting costly scheduling errors and reruns. By tapping the potential productivity of your CPU. By freeing your operators for more important tasks. And by making certain your data center runs smoothly, day in and day out.

To find out more, call Altai Software today. Once you see Zeke work, you'll agree that Zeke does a whole lot more. For a whole lot less.

ZEKE
The Scheduler that Works.



Software that works.

ALTAI Software • 624 Six Flags Drive • Suite 150 • Arlington, Texas 76011
800/227-7774 (answered 24 hours a day)

- diagmsg(treatment, diagnosis messages).
- helpmsg(trait, help messages).

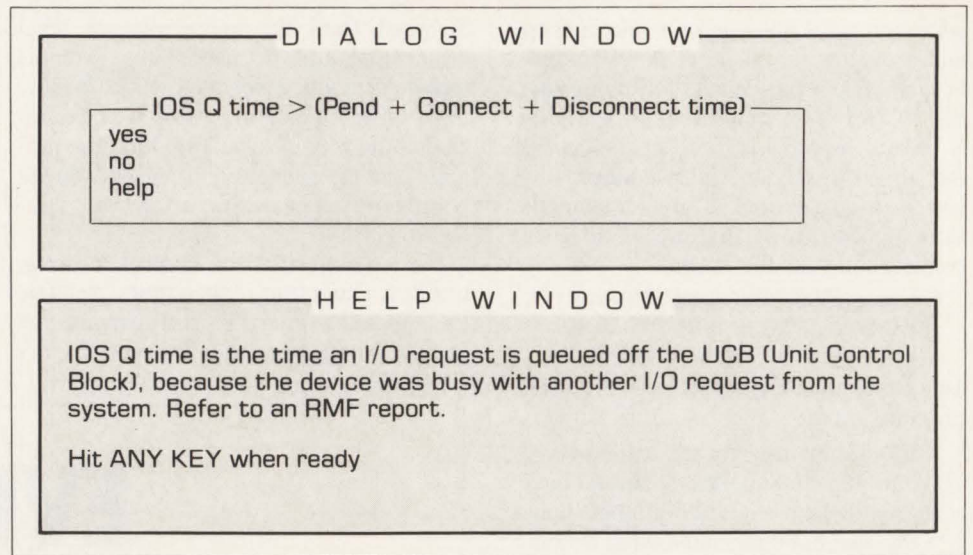
where:

- prompt associates with each trait the corresponding prompt message to ask the user. "Is" is printed before each prompt message implying a true/false question.
- condition associates a list of traits with a particular symptom. For example, condition(ratehi, [i1, i2, i3]) would imply that traits i1, i2 and i3 are all associated with the symptom ratehi. This knowledge is used to control the flow of prompting.
- depends will ask the prompt message associated with the specified trait only if the logical expression associated with the list of traits evaluates to true. For example, if the question, "Is VIO being used?" was answered positively, then it would make sense to ask, "Is VIO being used for small temporary datasets?" With the depends primitive, if VIO were not being used the "small temporary dataset" question would not be asked. Depends provides a finer level of control over the dialog that the condition primitive does not provide.
- diagnosis associates the true/false values of traits with whether a particular symptom exists. If the symptom does exist, treatment is the index to invoking the correct diagnosis and treatment messages. For example, consider:
 - diagnosis(ratehi, [i1, i3], fixrate).
 - prompt(i1, "Demand Paging High").
 - prompt(i3, "Storage Isolation Used for IMS or CICS").


This implies that the condition ratehi is present if traits i1 and i3 are both true: demand paging is high and storage isolation is being used for either IMS or CICS. A minus sign " - " in front of a trait, that is -i2, would imply that trait i2 is tested for falsehood. On closer examination, diagnosis represents how tuning rules are coded. If it is possible for a symptom to be caused several different ways, then each possibility can be coded. The inference mechanism automatically checks for all possible solutions:

- treatmsg, diagmsg - associate a treatment with a set of messages: treatmsg deals with treatments, diagmsg deals with diagnoses

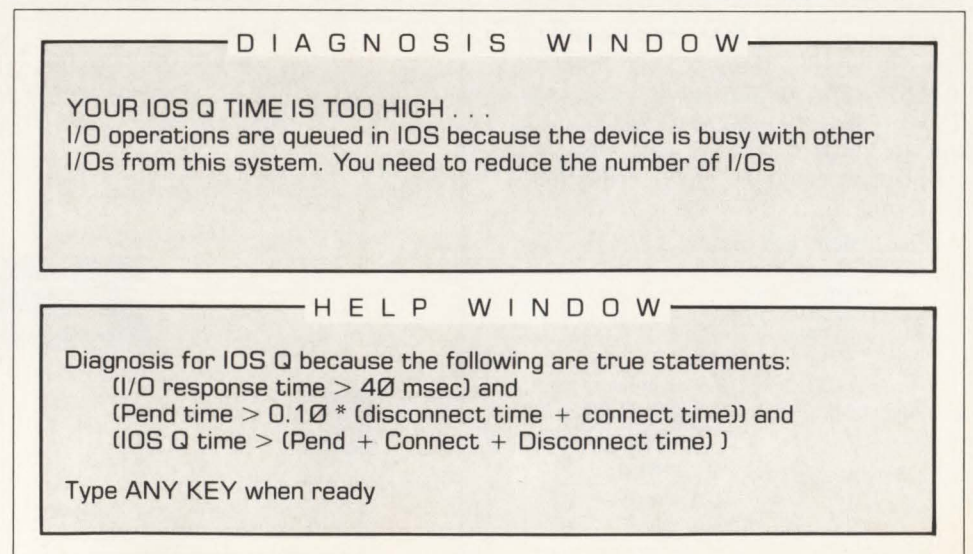
WINDOW 1



WINDOW 2



WINDOW 3



- **helpmsg** – associates a trait with a set of help messages that try to give more information about the particular trait.

Thus, each of the above represents knowledge and is kept in a file that can be edited by the user. The inferencing mechanism is thus reduced to processing these predicates: it is short and contains no knowledge. Conceptually, it can be used in any expert application as long as the syntax outlined above is followed.

Several features were added to the general inferencing mechanism. After a consultation, the user can invoke a built-in editor that resembles the PC word processing program WORDSTAR™. Rules can be added, deleted or modified. Care should be taken so as not to violate the syntax outlined above: no syntax checking is done. In addition, a “Why?” feature was added. After each diagnosis is given, the user is given the opportunity to ask the expert system why that diagnosis is being given; in other words, how that conclusion was reached. The inferencing mechanism backtracks from the treatment using the diagnosis and prompt predicates to replay the rules that were invoked using English.

A Consultation Session

It is difficult for the reader to get a firm understanding of the expert system described unless (s)he has the actual floppy disk with the program available. To eliminate this difficulty, this section will display parts of a sample consultation session with the expert system. The discussion will show the screen displays a user would see and will provide an overview of the overall flow of the system. Finally, a short discussion of how a knowledge base can be edited will be

Required Files

First, the program (henceforth called ADVISOR) must be present along with at least one knowledge base. Knowledge bases are ordinary text files that can be created by any ordinary PC editor that does not use embedded control characters. ADVISOR knows that files are knowledge bases if they end with the suffix KBA (for Knowledge BAse). In our prototype we created two knowledge bases: IO-TUNING and PAGING.

The user uses the PC arrow keys to select which knowledge base they want to use; in essence, this is actually equivalent to selecting a knowledge domain to be exploited.

At this time, we point out the different windows that ADVISOR uses:

- **DIALOG**: this window is the primary means of communication between ADVISOR and the user. Performance questions are posed to the user and each reply should be followed by hitting the RETURN key.
- **HELP**: should a user need assistance in understanding a question, the HELP window will display that information. HELP is invoked by replying with a ? to the question posed in the DIALOG window. In addition, the HELP window is used later to explain why a particular diagnosis was arrived at.

- **DIAGNOSIS**: after a set of related performance questions is asked, any performance problem found will be noted in the DIAGNOSIS window. Here, the user is given the opportunity of asking why a particular DIAGNOSIS is being made.
- **TREATMENT**: often, a particular diagnosis will have one or more possible remedies or TREATMENTS. Each possible TREATMENT is displayed in this window successively, each one being preceded by giving the user the opportunity to ask WHY.
- **EDIT**: once a consultation session has ended, the user is given the opportunity to edit a knowledge base. Here,

W I N D O W 4

— D I A G N O S I S W I N D O W —

YOUR IOS Q TIME IS TOO HIGH . . .
I/O operations are queued in IOS because the device is busy with other I/Os from this system. You need to reduce the number of I/Os

— T R E A T M E N T W I N D O W —

I/Os can be reduced by . . .
– increase block sizes and/or the number of buffers.

Hit ANY KEY when ready

W I N D O W 5

Line 1	Col 1	Indent	Insert
			prompt("resp","I/O response time > 40 msec")
			prompt("pend","Pend time > 0.10 * (disconnect time + connect time)")
			prompt("chan","any channel utilization at least 30% ")
			prompt("dpr","Dynamic Path Reconnect (DPR) supported on your DASD ")
			prompt("wkld","your primary workload batch")
			prompt("iosq","IOS Q time > (Pend + Connect + Disconnect time) ")
			prompt(" avsk","average seek distance between 40 and 50 cylinders")
			prompt("vtoc","cylinder boundary allocation used for VTOCs ")
			prompt("svio","VIO being used for small temporary datasets ")
			prompt("wvio","VIO being used for sort work files ")
			prompt("cvio","cylinder boundary allocation being done for VIO datasets")
			prompt("seq","the blocksize for sequential datasets between 4K & 9K ")
			prompt("part","the blocksize for partitioned datasets between 4K & 8K ")
			prompt("direct","the blocksize for direct datasets between 2K and 4K")
			prompt("compat","the blocksize for compatibility datasets = 6K ")
			prompt("pgm","the blocksize for program libraries = 18K ")
			prompt("tape","the blocksize for tape datasets at least 4K ")
			prompt("t6250","the blocksize for 6250 tape datasets = 32K ")
			prompt("other","the blocksize for 800/1600 tape datasets between 8K & 12K ")
			condition("pend",["resp","pend","chan","dpr","wkld"])
			condition("iosq",["iosq","svio","wvio","cvio"])

a full-screen editor is invoked that allows the rules in the knowledge base to be modified, deleted or extended.

Examine Window 1 to become familiar with the "look and feel" of the system. As each question is asked in the DIALOG window, the user can ask for HELP. When finished with HELP, typing any key will return the user to the DIALOG window where (s)he is reminded to reply with a *y* (yes) or an *n* (no).

Once all of the questions or traits corresponding to a logical sequence have been asked, descriptions of any performance problems detected appear in the DIAGNOSIS window.

In Window 2, the user is given the opportunity to ask what caused (WHY) this diagnosis to be made by ADVISOR. If the user replies *n* to the WHY prompt, the TREATMENT window appears (more later). If the user replies *y* to the WHY prompt, the HELP window reappears containing an English-like replay of the prompts used to arrive at the current diagnosis (see Window 3). Notice, if a particular prompt is preceded with a hyphen (-), this implies the negation of the particular prompt.

In Window 4, the user types any key which causes the HELP window to be replaced by the TREATMENT window. Here, a possible suggestion is made that the user can adopt to ultimately help tune the actual MVS system. As before, the user types any key which causes ADVISOR to continue.

If a particular diagnosis reappears following the disappearance of the TREATMENT window, it implies that the current diagnosis has been reached by invoking additional rules. In this case, the user is given a different remedy that might apply to the actual MVS system.

Once the consultation session has ended, the user is given the opportunity to edit a knowledge base.

As at the start of program execution, the user selects a knowledge base with the PC arrow keys followed by enter. Here, a full-screen editor similar to WORDSTAR™ is invoked that operates on the knowledge base (see Window 5). Finally, the session ends.

Issues for Future Development

As this PC/PROLOG based system is still in its infancy, several issues regarding where the system should next go need to be expressed. The following issues should apply to any expert system for CPE.

Location Help

Rather than just explaining the meaning of traits in help messages, add information about where to find measurement(s) associated with the trait; that is, on page 2, bottom of an RMF report. This is an example of using an additional class of knowledge that, until now, only the expert has.

Automatic Data Reduction

Ideally, a "front-end" mechanism that could read and analyze measurement sources like RMF automatically would greatly simplify the user's interaction with the system. Rather than putting the burden of responsibility on the user to get the answers to the prompts, a program(s) could be designed to extract the corresponding answers and download these to the PC. A potential source for much of this data could be Morino Associates' (Vienna, VA) MICS database which is a repository for many different types of measurements. The MICS data dictionary already assigns names to many measurements and these names could be related to the traits we referred to previously.

Links to Other PC Packages

If data were downloaded from the mainframe to the PC, it could be stored in a spreadsheet or a database. Products like 1-2-3 from Lotus Development or dBASE-III from Ashton-Tate could then be used to do the analysis on the PC rather than the mainframe. These could be invoked from PROLOG directly as the interface is already defined.

Certainty Values

Some rules are not as clear cut as one would like. If a certainty value, say between 0 and 100, was attached to each rule's conclusion, a certainty threshold could be specified for the premises. Below this threshold, rules are considered to be false. This would allow fuzzy reasoning to be a part of the system.

External Program Execution

As in many mainframe offerings, it is sometimes convenient to have an exit facility. This would allow the PROLOG system to escape to execute some user-specified program and then return. This could greatly increase the potential applicability of the system.

Conclusions

The objective of this article was to provide new insight into CPE by applying AI techniques. An expert system is a complicated object with many advantages. Foremost among these is that expert

knowledge is elicited and stored in an executable form.

The system described is only a beginning at applying current technology. It was never intended to replace the human expert. On the contrary, the original intent was to use the system as a training mechanism in which a new CPE practitioner could pose performance scenarios to the expert system and see how these would be analyzed and what recommendations would be made.

It is believed that similar applications will flourish in the next decade. Yet, though the size of knowledge bases may increase, the human inference mechanism is still better than any that will be built during that time. The article concludes with the following quote from Aristotle: "It is the mark of an instructed mind to rest satisfied with the degree of precision which the nature of the subject admits, and not to seek exactness when only an approximation of the truth is possible."

Acknowledgements

The author would like to take this opportunity to thank those referees and editorial review board members of the Computer Measurement Group for their constructive comments and suggestions. In addition, the contribution of Sid Soberman to the actual development of the system was critical to its eventual success. His comments and criticisms warrant thanks here. And though the work was conceived and performed solely by the author, he would like to thank those staff members of Bell Communications Research for reviewing the work as it progressed. Their input is gratefully appreciated here. ☺

ABOUT THE AUTHOR

Dr. Bernard Domanski is an associate professor at City University of New York where he earned his PhD in Computer Science and he is also an independent consultant. During the last 15 years, he has participated in the design, implementation and execution of benchmarks for IBM and IBM code compatible processors spanning the 370/168 through the IBM 3090-600E. He has written more than 30 papers and lectured internationally on subjects such as workload characterization, MVS performance and tuning and IMS. He has been the technical program chairman for the Computer Measurement Group since 1983 and has made numerous presentations at CMG, SHARE, GUIDE, CPEUG and ECOMA.

Selecting CASE Technology And Software

By Howard W. Miller

By definition, CASE (Computer Aided Software Engineering) is a technology that applies an automated engineering-like discipline to the phases of computer software system development, project management and software design, development, testing and maintenance. It is a highly flexible software development environment.

CASE is an opportunity to accelerate the expansion of information technology by resolving some of the software development roadblocks. It has the ability to improve the flexibility of the software development process and reduce the elapsed time and the effort associated with developing, maintaining and altering software systems.

The acronym CASE is relatively new, but some of the products that are classified as CASE tools have been around for a long time: code generators, design aids and project management tools. Other products are so new that they have not been subject to the test of time. Also, they are changing and evolving rapidly.

There are more than 50 products available that can be classified as CASE technology that are listed in the side bar to this article. Generally these products fall into two categories: toolkits and workbenches. A toolkit is a set of tools that supports one or more of the software development functions: planning, analysis, programming, maintenance and project management. Integration differentiates a workbench from a toolkit. A workbench is an integrated development environment

that supports either a full range of the software development life cycle or a major portion of the development cycle — the design or code generation portion.

There is no free lunch with CASE technology. The combination of new and old products, the range of products from toolkit to workbench and the sheer number of entries in the market make selection of CASE technology difficult. For some, the easiest solution to this dilemma is to do nothing. For others, an early investment will give them a competitive position. This article clarifies the CASE selection and implementation process.

Selecting CASE Technology

Clarify Objectives

Which benefits will be derived from the CASE tool: design, code generation, maintenance or project management? Most commercial organizations purchase a large portion of their software. However, when software is developed, it is developed with a fourth-generation language. It is not likely under these circumstances that a back-end workbench or a programmer toolkit would have the highest return. A front-end workbench, an analysis toolkit or even a project management toolkit would make more sense. Regardless, it is essential that you clarify your objectives from the start.

Narrow Options

Choose a CASE product category and analyze the tools within that category. Choose a tool that complements your or-

ganization's methodology or methodologies. If you do not have a methodology, *stop* your selection process and begin the process of selecting a methodology. When you begin the CASE product selection, consider the following:

- Hardware requirements — The options are mainframe, work station or PC-based toolkits or workbenches.
- Tools already in place — An existing DBMS, Data Dictionary or fourth-generation language increases the need for selecting complementary CASE product.
- Types of systems to be developed and supported — What are your software development strategies? The requirements for installing on-line mini-computer systems are different than on-line mainframe systems. Requirements for purchased software are different than for in-house developed software.
- Interviews with current users of product — Would the current user choose the same product again if given the chance?
- Product demonstration and hands-on experience — Do not overlook the opportunity to see the product in action.
- Available product support and training — Are support and training available? Will the vendor be in business next year or five years from now? Does the vendor have a good reputation?
- Your organization's future direction and that of the vendor — Most CASE

products are incomplete and in a state of transition. Is the stated direction of the vendor in concert with the direction of your organization?

Make your selection using an objective criteria. Choose on the basis of how well the toolkit or workbench does the following:

- Supports your methodology
- Integrates with other tools, DBMS, Data Dictionary and 4GL
- Integrates with other tools
- Supports diagramming
- Supports code generation
- Performs checking and analysis.

Implementing CASE Technology Methodology Selection

The installation of a development methodology is a prerequisite to CASE technology. CASE technology provides an automated engineering-like discipline for specification of computer system design, software development, testing and maintenance and project management. CASE, therefore, is automating the development process, the latter being the selected and defined methodology. Therefore, as a prerequisite, CASE requires that some methodology be in place for CASE to automate. There must be a management commitment to the development methodology. The staff and users must be trained and some level of experience is required.

It is very difficult, if not impossible, to embrace the concepts of a methodology and CASE technology simultaneously. There is an old adage of systems development that states if something cannot be done manually, it cannot be automated. This adage applies to methodology and CASE. Simultaneous installation of both a development methodology and CASE technology significantly reduces the chances of success for either.

CASE Selection

Select the CASE tools that support the methodology. Based on the organization, these could be support for single or multiple methodologies. Multiple division organizations or organizations that are likely to merge with others should seek flexibility through the use of tools that support multiple methodologies or that are methodology independent. Such organizations may find it advantageous not to impose standards that force an organization into a conversion. Smaller organizations can probably ignore this aspect of the selection process.

IBM Mainframe CASE Products

Toolkits

Planning Toolkit

Project Workbench
Applied Business Technology
365 Broadway
New York, NY 10013
212-219-8945

PC Prism
Index Technology Corporation
One Main Street
Cambridge, MA 02142
617-494-8200

Analysis Toolkit

DESIGN/1
Arthur Andersen
33 West Monroe Avenue
Chicago, IL 60602
312-580-0033

Excelsior/RTS
Index Technology Corporation
One Main Street
Cambridge, MA 02142
617-494-8200

ManagerVIEW
Manager Software Products, Inc.
131 Hartwell Avenue
Lexington, MA 02173-3126
617-863-5800

AION Development System
Aion Corporation
101 University Avenue
Palo Alto, CA 94301
415-328-9595

The Developer
ASYST Technologies, Inc.
1080 Beaver Hall Hill
Suite 1400
Montreal, Que, Canada H2Z 1S8
800-361-3673/514-871-0108

IDMS/ARCHITECT
Cullinet Software
400 Blue Hill Drive
Westwood, MA 02090
617-329-7700

EXELERATOR
Index Technology Corporation
One Main Street
Cambridge, MA 02142
617-494-8200

AUTO-MATE PLUS
Learmonth & Burchett
Management Systems-LBMS, Inc.
2900 North Loop West, Suite 800
Houston, TX 77092
713-682-8530

ProKit*ANALYST
McDonnell Douglas, ISG
P.O. Box 516, Dept. L863
Bldg. 302, 1E
St. Louis, MO 63166
314-232-5715

ProKit*WORKBENCH
McDonnell Douglas, ISG
P.O. Box 516, Dept. L863
Bldg. 302, 1E
St. Louis, MO 63166
314-232-5715

SYSTEM ARCHITECT
Popkin Software Systems, Inc.
60 Skiff St., Suite 251
Hamden, CT 06517
203-248-9700

Analyst/Designer Toolkit
Yourdon, Inc.
1501 Broadway, Suite 601
New York, NY 10036
212-391-2828

DesignAid
Nastec Corporation
24681 Northwestern Highway
Southfield, MI 48075
313-353-3300

VIA/INSIGHT
VIASOFT, Inc.
3033 North 44th Street
Phoenix, AZ 85018
602-952-0050

Data Analysis Toolkit

Data Resource LEVERAGE with JANUS
D. Appleton Company, Inc. (DACOM)
1334 Park View Avenue, 220
Manhattan Beach, CA 90266
213-546-7575

DESIGNMANAGER
Manager Software Products, Inc.
131 Hartwell Avenue
Lexington, MA 02173-3126
617-863-5800

SQL*Design Dictionary
Oracle Corporation
20 Davis Drive
Belmont, CA 94002
800-345-DBMS

Programming Toolkit

POLYMAKE
POLYTRON Corporation
1815 NW 169th Place
Suite 2110
Beaverton, OR 97006
503-645-1150

P-TOOLS
Phoenix Technologies, Ltd.
320 Norwood Park South
Norwood, MA 02062
617-769-7020

POLYTRON VERSION CONTROL SYSTEM (PVCS)
POLYTRON Corporation
1815 NW 169th Place
Suite 2110
Beaverton, OR 97006
503-645-1150

Maintenance Toolkit

ADPAC CASE TOOLS
Adpac Corporation
340 Brannan St.
Suite 501
San Francisco, CA 94107
415-974-6699

CASE Pilot

If you are choosing a workbench, the best route is to select a real but non-critical project to implement the CASE tools. Train your best and most receptive staff on the CASE tools. If faced with the decision between the best and the most receptive, choose the most receptive. After training is complete, put together a plan identifying the expected results and the measurement criteria. Implement the project and evaluate the results based on the evaluation requirements. Do not expect high productivity, because a large portion of the productivity gains from CASE technology comes from the reusable aspect of the CASE design elements.

However, if you choose a toolkit, you may choose to implement it across the entire department immediately. The procedural impact of installing an analysis toolkit, for example, will not be nearly as traumatic as installing a front-end workbench. It is, therefore, possible to implement it quickly and realize the benefits of the toolkit immediately.

CASE Implementation

After the pilot is complete, formalize the CASE procedures and make CASE the standard for all development projects.

Cautions

Three aspects of CASE require some words of caution. First, CASE is an emerging technology. Although there are more than 50 CASE tools identified in the side bar accompanying this article and even more available, none meet all of the requirements of the ideal CASE model. Recognize this shortfall at the start and select tools that maximize the benefits in your environment. Look for vendors and/or strategies that show potential for providing long-term solutions to the inadequacies.

Second, many CASE tools are tied to specific structured design methodologies, databases (IMS, IDMS, ADABAS), programming languages (COBOL, SQL) and hardware vendors. The marketplace is being driven by what is marketable, not necessarily what is best. Software vendors make a large financial commitment to developing their product. To achieve a financial return, they target the largest markets or the most abundant technology.

Select products that are flexible and products that support multiple design methodologies, databases, programming

IBM Mainframe CASE Products Continued

RECODER

Language Technology Inc.
27 Congress Street
Salem, MA 01970
617-741-1507/800-RECODER

INSPECTOR

Language Technology Inc.
27 Congress Street
Salem, MA 01970
617-741-1507/800-RECODER

Reverse Engineering

Meta Systems
315 East Eisenhower Pkwy.
Suite 200
Ann Arbor, MI 48106
313-663-6027

GAMMA

KnowledgeWare, Inc.
3340 Peachtree Road NE
Suite 2900
Atlanta, GA 30026
404-231-8575

SOURCEMANAGER

Manager Software Products, Inc.
131 Hartwell Avenue
Lexington, MA 02173-3126
617-863-5800

TELON

Pansophic Systems, Inc.
709 Enterprise Drive
Oak Brook, IL 60521
312-572-6000

TRANSFORM

Transform Logic Corporation
8502 East Via de Ventura
Scottsdale, AZ 85258
602-948-2600

NETRON/CAP

Netron, Inc.
99 St. Regis Crescent North
Downsview, Ont, Canada M3J 1Y9
416-636-8333

APS Development Center

Sage Software, Inc.
3200 Monroe Street
Rockville, MD 20852
301-230-3200

Integrated Workbenches

PACBASE

CGI Systems, Inc.
1 Blue Hill Drive Plaza
Pearl River, NY 10965
914-735-5030

MAESTRO

Softlab, Inc.
188 The Embarcadero
Bayside Plaza, 7th Floor
San Francisco, CA 94105
415-957-9175

EXSYS

DATA ARCHITECTS, INC.
245 Winter Street
Waltham, MA 02154
617-890-7730

CONTROLMANAGER

DATAMANAGER

DICTIONARYMANAGER

METHODMANAGER

Manager Software Products, Inc.
131 Hartwell Avenue
Lexington, MA 02173-3126
617-863-5800

Education

Digital Consulting
Six Windsor Street
Andover, MA 01810
617-470-3880

CASE EDUCATION OFFERINGS

- Computer-Aided Software Engineering Symposium
- CASE Symposium for Aerospace, Defense and Engineering
- CASE Benchmarks: A Seminar Comparing 7 Leading CASE Tools
- Software Engineering and CASE Technology
- Application Prototyping: Implementing The New System Development Technology

Workbenches

Front-end Workbenches

FOUNDATION

Arthur Andersen
33 West Monroe
Chicago, IL 60602
312-580-0033

INFORMATION ENGINEERING WORKBENCH

KnowledgeWare, Inc.
3340 Peachtree Road NE
Suite 2900
Atlanta, GA 30026
404-231-8575

EPOS

SPS
14 East 38th St.
14th Floor
New York, NY 10016
212-686-3790

vsDesigner

Visual Software, Inc.
3945 Freedom Circle
Santa Clara, CA 95954
408-988-7575

INFORMATION ENGINEERING FACILITY

Texas Instruments
P.O. Box 655621 M.S. 8474
Dallas, TX 75265
214-575-4404

Back-end Workbenches

Life-cycle Productivity System

American Management Systems
1777 North Kent Street
Arlington, VA 22209
703-841-6000

AutoCode

Integrated Systems, Inc.
2500 Mission College Blvd.
Santa Clara, CA 95054
408-980-1500

CORVET

Analysts International Corp.
354 Eisenhower Pkwy.
Livingston, NJ 07039
201-535-9844

CorVision

Cortex Corporation
138 Technology Drive
Waltham, MA 02154
617-894-7000

languages and hardwares. The nature of business is such that mergers and acquisitions are common. Further, technology is changing. There may be a desire to support a relational database with SQL on hardware from different vendors. Maintain flexibility wherever possible.

Lastly, CASE productivity gains come from two areas: reusability and maintenance. CASE productivity comes from developing an extensive library of design specifications that can be cut and pasted to create new applications. The benefits of reusable designs are not immediate. It takes time to develop these inventories; therefore, benefits are deferred. The second area of productivity comes from making software reliable at the start and maintainable as time goes on. This reduces the spiraling expense of software maintenance. Again, the benefits are deferred and only occur for software that is designed with CASE or if an initial investment is made in reverse engineering.

Conclusions

If there is one major fault with the information technology industry, it would

be the constant search for heroic solutions. The profession has embraced the use of such diverse technology as third-generation languages (COBOL, PL1), fourth-generation languages, structured methodologies, database, personal computing and so on in hopes that all problems would be solved. Now, the solution is CASE technology.

CASE is real and it is usable. There are between 50 and 100 products on the market that qualify as CASE tools. Some sets or subsets of these tools will benefit most software development environments. However, CASE is not a panacea. Only when CASE is orchestrated with other techniques and technologies such as sound business planning techniques, unattended operations, artificial intelligence, personal computing and cooperative processing will CASE achieve its highest return.

CASE is also new and incomplete. Many inconsistencies and annoyances can often be fixed in the next release of the product. Many CASE vendors are coming out with new releases two or more times a year. When selecting a CASE product, take these factors into consideration. If such a volatile environment is not ac-

ceptable, you may not want to get into CASE at this time.

CASE addresses the issue of quality. Studies have indicated that less than 25 percent of all software projects come in on time and within budget and that less than 25 percent of all projects meet the expectations of the requestor. Therefore, less than six percent of all projects satisfy the needs of the requestor in the time frame that they were requested. CASE addresses this problem. It automates the software engineering process, improving quality. By reusing the automated design elements and through reverse engineering, CASE enhances productivity and provides a *high quality system*. ☉



ABOUT THE AUTHOR

Howard W. Miller, CDP, CSP, is responsible for administrative computing at Boston University. He has held senior-level positions in system management for more than 20 years. He is characterized as a businessman whose area of expertise is business information technology.

He holds a BA degree from Duquesne University (Pittsburgh, PA) and an MS degree from the University of Bridgeport (Bridgeport, CT).

Full Lifecycle CASE Methodology

Data-driven design – Dictionary-driven development

MAGEC is the first active dictionary to automate application design, development, control, and maintenance. Greatly improves both productivity and quality of applications for both on-line and batch systems.

- Prototyping
- Report Painting
- Documentation
- Debugging Aids
- Screen Painting
- Program Generation
- Security
- PC Workstation

MAGEC operates in both mainframe and PC environments, it supports MVS, VSE, VM/CMS, MS-DOS, and OS/2. Applications can be ported with **no programming changes** required.

MAGEC® – For professionals serious about productivity improvement

MAGEC Software™
P.O. Box 260319
Plano, TX 75026-0319

(800) DD MAGEC
(800) 336-2432
(214) 248-0823

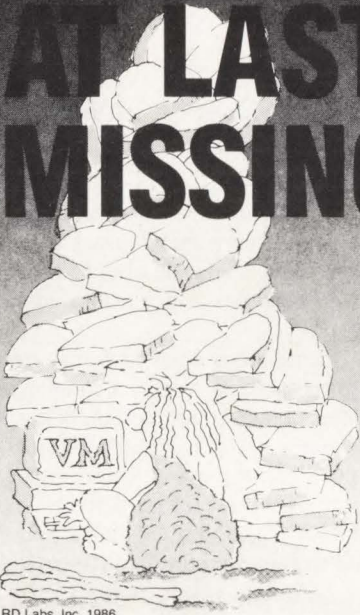
For information on advertised products, fill out the
READER SERVICE CARD
elsewhere in this magazine.

ADVERTISER INDEX	PAGE
ACTS CORPORATION	73
ALTAI SOFTWARE, INC.	58, 59
B I MOYLE ASSOCIATES	82
BENNETT SOFTWARE, INC.	18
BLUELINE SOFTWARE, INC.	13
BMC SOFTWARE	55
BOOLE & BABBAGE	51
CANDLE CORPORATION	
..... 16-17, 48-49,	93
CHANEY SYSTEMS SUPPORT	34
COMPUWARE CORPORATION	47
CYBERMATION INC.	71
DATA 21	92
DAVIS, THOMAS & ASSOCIATES, INC.	70
DOWLING ASSOCIATES	83
EMC CORPORATION	7, 87
GENERIC SYSTEMS	78
GOAL SYSTEMS INTERNATIONAL	96
H & W COMPUTER SYSTEMS INTERNATIONAL	45, 84
HAMILTON COMPUTER SALES & RENTAL	80
INNOVATION DATA PROCESSING	77
INTEGRITY SOLUTIONS	43
INTERNATIONAL SOFTWARE TECHNOLOGY	92
IPL SYSTEMS	79
LANDMARK SYSTEMS CORPORATION	3, 5
LAWSON ASSOCIATES	21
LEVI, RAY & SHOUP, INC.	53
MACKINNEY SYSTEMS	88
MAGEC SOFTWARE (AL LEE & ASSOCIATES)	66
MHT SERVICES, INC.	82
NBS SOUTHERN, INC.	57
NITTANY SYSTEMS, INC.	72
PRINTER SYSTEMS CORPORATION	39
QUANTUM INTERNATIONAL CORPORATION	81, 83, 85
RD LABS, INC.	67
SDI	41
SMARTECH SYSTEMS	37
SOFTOUCH SYSTEMS	71
SOFTSYSTEMS, INC.	20
SOFTWARE PURSUITS	67
SOFTWARE WORKS, INC.	46
SORBUS	11
SYSTEMS & TELECOMS	27
UNIVERSAL SOFTWARE, INC.	85
VM1	89
VM SOFTWARE INC.	2
WESTINGHOUSE MANAGEMENT SYSTEMS	9

If you have ideas, comments or suggestions concerning MAINFRAME JOURNAL, write or call: Bob Thomas, editor-in-chief, MAINFRAME JOURNAL magazine, PO Box 38185, Dallas, TX 75238, (214) 343-3717.

NEW RELEASE!

AT LAST, THE MISSING LINK



RD Labs. Inc. 1986

RD/SHARE: VM/CMS

SOURCE CODE LIBRARY MANAGEMENT AND FILE SHARING.

VM/CMS has XEDIT and UPDATE for file maintenance. But you have not had a library tool for storage, sharing, auditability and security of your source files—until now.

- Source code management with version control and history
- Comprehensive management reporting capability
- Flexible security system with a full audit trail
- All this and more: Most systems **less than \$10,000**

STEP OUT OF THE STONE AGE WITH RD/SHARE®

Call for more information or a 30 day trial:
1-800-346-5227
IN CANADA & CALIFORNIA: 916-624-5755

RD LABS
THE VM SPECIALISTS

CIRCLE #68 on Reader Service Card

Break through the Report Distribution Bottleneck. Exchange Bank did. With SPRI/EXPRESS.

Exchange Bank, a ninety-year old, 16 branch financial institution, had a problem. 200 computer reports going in 16 different directions. All of which had to be manually sorted and collated for distribution.


Until **SPRI/EXPRESS**—Software Pursuits' powerful report distribution system for VSE users. Now, all reports destined for a single user are printed together in a customized package ready for distribution. No more long hours spent collating. No more lost and misrouted reports.

Best of all, **SPRI/EXPRESS** uses a fraction of the disk space and processing time required by the other report distribution systems Exchange Bank evaluated.

Automate your report distribution effort with **SPRI/EXPRESS**. Call toll-free for more information (800) 367-4823, in California (800) 367-9851.

SOFTWARE PURSUITS 
The proven alternative for VSE users.

1420 Harbor Bay Parkway, Suite 200, Alameda, CA 94501



"SPRI/EXPRESS has already saved us almost one full-time position."

—Gary Searby,
Data Center Manager,
Exchange Bank,
Santa Rosa, CA

CIRCLE #4 on Reader Service Card

Apply Large Network Management Strategies To Smaller Networks

By Stephen N. Steinmetz

Today's networks are critical and expanding. Companies with large data communications networks have learned how vital their networks are. So have companies who use their networks extensively for customer service, sales (that is, airlines, car rental companies, stock brokerages and so on) and compliance with government regulations for keeping extensive, current records (that is, financial institutions, medical providers and so on). A network outage in one of these companies causes a definite negative impact on their business objectives that is easy for all to understand.

The organizational structures, tools and strategies that these companies use to effectively manage their networks may not be justifiable for a company with a smaller, less visible network.

However, many smaller companies have realized that they are just as dependent on their data communications networks for growth or just plain survival. Effective network management is just as vital to these companies.

If you know how critical your network is to your business and the degree of criticality does not indicate the need to incorporate "large" network management strategies now, the potential for network expansion may.

Networks always seem to expand. New applications are purchased or developed to serve a new or existing business requirement. Entry of company information into the computer is more and more the responsibility of the department that gathers it. Successful implementation of electronic mail requires computer access by nearly all of the staff. More staff members are introduced to the company's

computer systems. As more applications, terminals and dependent users are added to a network, the task of managing that network becomes much more vital and complex.

By adapting strategies now that are used to manage large networks, a company with a smaller network can provide better service to users of its current network. This also positions the company for managing and supporting its future network (which I maintain will be larger, faster and more crucial to the company's business).

Strategies for Large Networks

Companies with large or critical networks can cost justify large network management staffs and expensive network management tools. They also recognize the importance of coordinating and planning all network activity. The reason is usually that they have survived the problems and negative business impacts that result from haphazard network management.

Smaller companies can learn much from how large network management has evolved. A *full* treatment of network management concepts is beyond the scope of this article. However, this article will help you understand why and how to apply large network management strategies to smaller networks.

Network Management Concepts

Network importance recognized

Companies with critical networks have usually recognized the importance of their networks the hard way (after a major problem). Once the critical nature of the network is realized, many companies try to quantify the cost of network outages and slowdowns. One airline estimated

network outage losses at \$10,000 a minute. Knowing the potential costs of a network outage is essential to the network design and budgeting processes.

Network management mission defined

The mission of a network management/support organization is to provide users of a company's data communications network access to that company's (and in some cases, other companies') computers. In doing so the organization must meet the availability, response time and long-range network objectives determined by management.

Service levels defined

Network availability and response time requirements should be derived from the business needs of the company. These requirements are called service level requirements.

Like any other service that data processing provides to the company, network requirements are driven by the department that requests the service. The user department sets these service requirements based on business needs. In heavily computerized companies, the user department is charged for the cost of the requested service. Since these costs are included in that department's budget, cost justification is the responsibility of the user department.

There are hardware and software tools available to accurately measure the level of service that is provided to users. Companies that have these tools in place implement service level agreements. Service level agree-

ments are formal agreements between the data processing department and the user departments that specify the service levels that will be provided to the user for a set monthly charge. Failure to provide the agreed upon service level results in a lower charge to the user and a budget overrun for data processing that month.

Also, as business needs or costs change, these agreements are renegotiated and the monthly charge to the user group is adjusted accordingly.

Long-range communications goals defined

The cost of providing data communications services to staff members is not trivial. Unless this cost is weighed against the benefits provided, it often seems astronomical. One way to help minimize the costs and maximize the benefits is to include data communications in the company's strategic business planning process.

If long-range business goals are shared with the network design group, a long-range communications plan can be developed. Projects can be completed with the overall plan in mind and duplication of effort and waste can be avoided. For many companies, connectivity to different types of computers may be a requirement (current or future). If the long-range plan requires connectivity to be addressed, equipment evaluation is at least partially based on this issue and the usable life span of any new equipment is lengthened. In addition, sharing resources is maximized and implementation of inadequate, quickly obsolete systems is avoided.

Another potential for increased benefit comes from input of ideas that a network design group can contribute to the strategic planning process. Network designers follow new technology and may see something that could give the company a competitive edge or a new business opportunity.

Network management functions defined

Certain functions must be accomplished to perform this mission effectively. These functions include:

- Providing the users with network problem help, information and general usage assistance

Network Support Organization Evolution

Networks usually start with a few terminals here and there. Another application is added, then another, then another and with each new application a few more users are added until, suddenly, the network is so large that the whole company depends on it.

Network support organizations often do not grow as gradually as the network does. Companies that do not recognize the rate at which their networks are growing tend to try to "make do" with the people and tools that have been in place until network outages become visible business problems. When that happens, there is a flurry of activity, a hasty reorganization and a full-scale network support organization is formed.

This is a painful process for all concerned. The users and customers suffer and complain until the problem is recognized and corrected. The managers who are responsible for the network have to scramble and the new network support staff has to cope with a tremendous learning curve. Many extra hours and dollars are spent.

It does not have to be that way!

If it is recognized that the network is, or will be, crucial to accomplishing the company's goals and if the proper plans are made, supporting a growing network can be smooth and painless.

Assuming there is recognition of how critical the network is now, or will be in the future, the next step is to plan for a support organization that will be sufficient for the future. Once this is done, the only problem is how to get from here to there cost effectively.

To use a highly technical networking term, this is a "piece of cake."

It starts by assessing the technical skills (and desires to learn networking) of current staff members and adding the network support tasks to their job descriptions. If the network is very small, the staff should be able to keep up with it as they probably have been doing the work already. All you are doing is formally recognizing that these tasks must be done.

This may cost the company a little at review time, but it ensures that the work is addressed on a regular basis. It starts instilling a "pride of ownership" feeling in the affected staff members and provides some incentive for staff members to develop network support skills. It also helps justify formal education for staff members. Education lays a strong technical networking foundation, brings in fresh networking ideas and boosts staff morale.

Next, institute a regular network review meeting (weekly or bimonthly is good). Use this meeting to review network performance, coordinate network projects and monitor growth. During each meeting have a staff member give a short review of new network ideas or technology on a rotating basis.

As the workload changes, make appropriate changes to the organization. Add new positions and either move current staff members around or hire new people or both. For example, in a company with a small network, the computer operations staff usually handles network operation. As the network grows, the prime-time operators have difficulty handling both computer operations and network operations tasks. As a result, there is no coordination of network operations tasks. By designating one computer operator as the primary network operator, most network tasks are done by this person. This accomplishes three things:

1. All network tasks are coordinated
2. The designated network operator learns more and more about network operations helping to improve service and build expertise
3. Someone in the organization has a finger on what the network is doing at all times.

When the amount of network tasks reaches a level where the network operator can no longer successfully accomplish the required computer operations tasks, it is time to create the network operator position. Then move that person into it and hire a new computer operator. Or hire a network operator from the outside.

A small (network-wise) company may have users calling the computer room to report problems. If the computer room staff cannot take the call, a secretary or answering service may get it. This leads to poor service and user frustration. Publish a number for any user to call with any computer-related problem and assign the job of answering that phone (and resolving or referring the problem) to a staff member. When the volume of calls interferes with the staff member's other responsibilities, create the help desk operator position and fill it.

The same strategy works well for the technical support, applications support, PC support and management areas.

Realizing the importance of the network and how large the network may become justifies the planning process. Once the plans are in place, it becomes easy to cost-effectively support an expanding network.

Stephen N. Steinmetz

- Monitoring the network for outages and response time
- Identifying and repairing (or circumventing) malfunctioning hardware
- Identifying and repairing (or circumventing) malfunctioning or improperly configured communications software
- Identifying and repairing (or circumventing) malfunctioning

- applications software
- Reporting network status and problems to management (this includes keeping good network and problem documentation)
- Measuring and reporting performance trends for input to the planning process.

A network management organization defined

In a company with a large or crit-

ical network, there may be one or more persons to accomplish each of the above tasks. In a company with a smaller network, one person may be responsible for several of these tasks. Organizationally, these people are usually divided into six functional areas. These areas and functions are:

Help desk — a single phone number for all users to call. Any type of computer-related problem is reported here, simple problems may be resolved, more complex problems are referred to the appropriate people and all problems are tracked to resolution. After a problem is resolved, it is the responsibility of the help desk to check with the user to ensure that resolution is complete. Network and problem status reporting is also performed here.

Network control center — where the next level of network problem determination and resolution is performed. Outages are investigated, isolated to the component level and resolved if possible. If the problem cannot be resolved, this group then circumvents the failing component if possible and refers the problem to the next level of support. The staff keeps statistics on response time and availability, assists the hardware group with installations and helps perform feasibility testing of new technology.

Technical support — those responsible for system software. They are the third level of support for any network system software problems. When the software vendor is required to fix a problem, it is the technical support staff that interfaces with them. These people are also responsible for installing, customizing, maintaining and helping to select system and system support software. They are also involved in the selection of computer and network hardware.

Applications support — this staff is the second level of support for any computer application-related problem. This includes program malfunctions and usage problems. These people are responsible for final testing of new applications and ensuring that users are trained and have adequate instructions in application use.

PC support — with the prolifera-

DTA[®]

Software Products

Effective Solutions for Technical Problems

DTA/RECOV - Feel secure knowing you can **recover** your critical online **VSAM datasets** in a matter of minutes. Forward/Backward recovery using standard CICS journals, even while CICS is running. Comprehensive reporting and a verification system provides confidence in the recovery process. **\$4800 for VSE, \$6000 for MVS.**

DTA/JOURNAL - Optional feature available with DTA/RECOV. Provides the same **Journaling Protection** for your **Batch Programs** that you already have with CICS. **NO recompilation** of your DOS/VSE and MVS applications. Reduces time consuming VSAM master file backups. Now you can have complete journaling of all updates to online or batch VSAM datasets. **\$2000 for VSE, \$2400 for MVS.**

DTA/PRINT - Provides simple and complete control of remote printing, viewing and management of your **VM** and **POWER** Queues from any **CICS** terminal. Computer Operators, Programmers, and end users love its ease of use. Priced at only **\$5000 for POWER**, only **\$7000 with VM** interface.

DTA/ICCF - Provides access to any CICS transaction without going through the Interactive Interface or using **/LOGOFF.../LOGON**. **\$2400**, or **\$800** when purchased with DTA/PRINT.

DTA/QCOPY - VSE utility that **Copies** or **Moves** reports from the **VM** or **POWER** Spool Queues to the **POWER** or **PUN** Queue. Complete or partial reports can be selected and combined into new reports. **\$3000**, or **\$2000** when purchased with DTA/PRINT.

DTA/COBOL - Give your COBOL programmers the **information** they need (where they need it) to perform the hardest task they have - **to maintain someone else's program**. Runs as part of your normal compile process. Reformats the standard COBOL output, reducing printed output by up to 50%, yet providing more information. **\$1800.**

DTA/DOC - View your VSE Hardcopy file from any CICS terminal. Same basic VSE commands plus many enhanced commands. Priced at only **\$800.**

DTA/LMON - Your systems programmer needs this tool to make efficient use of your **SVA**. Shows you how to return **wasted Virtual Storage** to more productive use, and how to reduce repetitive loads to your system libraries which will improve batch throughput. Priced at only **\$1700.**

Customer Support - All products are written and **supported** by the DTA **development staff**. When you need help, you talk with the experts on our toll free customer hotline.

For details on any of these effective solutions, call toll free 800-521-6773, in Minnesota, 612-591-6100.

DTA[®]

Davis, Thomas & Associates, Inc.
Software Products Group
550 Waterford Park
505 North County Road 18
Minneapolis, MN 55441


tion of personal computers, many companies have created a PC support function. This group is sometimes called the "information center" and is the second level of support for any personal computer-related problem. Additional duties include PC configuration and setup, PC inventory control, PC user training and evaluation of new PC technology.

Communications hardware — another separate group in some companies is this department. All voice and data communications facilities are evaluated, ordered and administered by this group. They are responsible for the installation of all communications hardware.

In many companies the voice facilities may be handled by an office management group and the data facilities by the technical support group. Creation of a separate communications department sometimes helps to best utilize the company's communications facilities, as long as some data communications expertise is included or available. Common voice/data facilities management is

ESP

EXECUTION SCHEDULING PROCESSOR™



The new standard in scheduling

ESP is the most versatile MVS scheduler available. The advanced design of ESP lets you manage your workload, whether you need a comprehensive system or want to phase in specific functions. Don't wait until problems get out of hand — ask us to demonstrate the flexibility of ESP today.



CYBERMATION INC.
16 ESNA PARK DRIVE
MARKHAM, ONTARIO
L3R 5X1
(416) 479-4611

Systems Excellence

CIRCLE #170 on Reader Service Card

CICS — WINDOWS™

Multiple Session on Terminals

- Save time.
- Have direct access to any application.
- Up to nine sessions on each terminal.
- Fast response with low resource usage.
- Easy to use.

Window Your Current Applications

- Instantly divide your current terminal screen into windows.
- Vary the size of the windows.
- Use your current applications within the windows.
- Save thousands of dollars in costly hardware changes.



SOFTOUCH SYSTEMS, INC.
8269 South Walker
Oklahoma City, Oklahoma 73139
405/632-4745

FOR MORE INFORMATION OR FREE TRIAL OFFER CALL OR WRITE TODAY 1-800-544-3036

COMPANY _____

NAME _____

TITLE POSITION _____

ADDRESS _____

TELEPHONE _____

CITY _____

STATE _____

ZIP _____

Automatic Data Compression

- Improve response time throughout your entire CICS network.
- Optimize line transmission speed, both local and remote.
- Save up to 60%.

CIRCLE #47 on Reader Service Card

essential if voice and data integration is to be implemented.

Network management tools implemented

There are many hardware and software network management tools available. Selection of tools is driven by service level requirements, network configuration and long-range plans. Each company's needs are as unique as its network.

Most companies with large networks have tools to do the following:

- Enter and simplify network control commands
- Monitor and report network performance for service level compliance and capacity planning
- Manage network problems and changes
- Test network facilities
- Repair or circumvent failing network hardware.

Network standards and procedures followed

Based on current needs and long-range plans, network standards and procedures are established and followed. Common standards include: naming conventions for hardware and software components, types of physical connections to be used and minimum documentation standards. Common procedures include: problem tracking and escalation procedures, component testing procedures and status, performance and problem reporting procedures.

Relating Strategies

These large network management strategies may be adapted to managing small- and medium-sized networks with tremendous success. The key to this is to plan for supporting a network larger than you

can imagine your company ever having. Talk to someone whose company has a large network and ask if he or she ever thought it would get anywhere near as large as it is today!

Planning for network expansion means that you start structuring your network support area now and grow into it. By following a few simple, inexpensive steps now, your staff can get better service from the current network (a positive business impact) and perhaps not even notice that the network is growing. Ten steps to follow are:

1. Evaluate how critical your network is to your business. (See "What Happens To Your Business When Your Network Goes Down?" in the January/February 1988 issue of MAINFRAME JOURNAL for a methodology.)
2. Review your long-range business plans now. Look for business impacts on the network and business opportunities from the network.
3. Develop a strategic communications plan. This should be a framework for all projects to follow. Allow for flexible (but coordinated) application of current and undreamed-of technology. Be sure to add network recovery to your corporate disaster recovery plan.
4. Define service level objectives for network performance to meet the business needs of each functional area in the company. As you acquire the means to accurately measure service levels, consider implementing formal agreements to ensure that service levels are based on business requirements.
5. Plan your future network support organization. Put together a tentative

organizational chart and prepare a mission statement.

6. Formally recognize the tasks that need to be performed to support your network. This means adding job responsibilities to some staff members' job descriptions. As the amount of network tasks to be performed exceeds the amount of time available, create the new positions and staff them (you will have developed some expertise inhouse). The accompanying side bar traces possible support organization evolution.
7. Create the help desk position as soon as possible and acquire (or develop) a problem management system. This will help track network availability and problem areas. All staff members will be able to use the computer more productively.
8. Review current network standards and procedures. Determine how much network expansion can occur before they become inadequate. Change standards and procedures to reflect any possible amount of growth. This is easier, safer and cheaper to do now than at the time they become inadequate.
9. Plan to acquire network management tools as they become necessary.
10. Be aware of changing network usage, requirements and technology. This should be an ongoing task until the day networks become obsolete.

Recognize how critical your current network is to your business and realize that your company will become more dependent on the network as it expands (which it will). It makes sense to model your network management strategies after those used by companies with very large or very critical networks.

These strategies help provide the best possible service to current and future network users (thus helping the company achieve its business goals) and allow for more orderly, less painful network expansion.

Implementing these strategies now, rather than waiting until it is essential, is also much less expensive in the long run. The cost of implementation is proportional to the size of the network.

I suggest that you start today. ●

ABOUT THE AUTHOR

Steve Steinmetz is a communications specialist with Davis, Thomas and Associates, Inc., a Minneapolis-based information systems consulting firm.

HUB/SDF

Screen Development Facility for CICS/BMS

- » Develop and generate CICS screens on mainframe and/or personal computer.
- » Converts existing BMS screens into HUB/SDF.
- » Develop screens on mainframe or PC with total portability between systems.
- » Supports COBOL naming conventions.
- » Supports repetitive columnar fields.
- » Prototype screen sequence processing with no application programming.
- » Review screens developed without generating BMS code.
- » Complete print documentation of screens developed.
- » Free demo disk available.

Nittany Systems, Inc.
(814) 466-6420

P.O. Box 452
Boalsburg, PA 16827

PC Version - \$350
Mainframe Version - \$1000

Manager *from page 50*

applications. In addition, you do not need to add several compression products to cover all your subsystems before you can realize appreciable reduction in network traffic.

Compression of nulls and spaces in 3270 data streams using the repeat to address order reduces network traffic by an average of 35 percent. On remote lines where the BAUD rate is the limiting factor, a 35 percent reduction in character transmission is significant. Besides the ability to suppress nulls and blanks, other compression techniques are also available. One technique is to compare the data stream being sent to the terminal with the session manager's image of the screen. The session manager would not send the data that is already displayed on the terminal screen. This would cut down the number of headings, names and static fields that would be retransmitted through the network to the terminals.

Another compression technique is to

process read buffer requests on the session manager level and not pass those read requests on to the terminal controller. The better session managers keep the same image of the terminal as the controller. Therefore, processing the read buffer request is a matter of emulating that task of the controller in the session manager.

There are other compression techniques and combinations that are available today, but the most common techniques are those discussed here. They result in the greatest reduction of traffic with a minimum of host overhead.

Conclusion

Multiple session management software is a cost effective set of tools that any company can afford. The boost in user productivity gained by providing access to several applications concurrently from a single terminal justifies the cost of most of these products. When you also consider the benefits of access security, help desk and network performance features, I cannot imagine working without a session manager at any installation. ☺



ABOUT THE AUTHOR

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

TEN GOOD REASONS...

to Let ACTS Corporation Handle Your Technical Training

1. Technical Awareness for Managers.

4 one day seminars: Systems Basics, MVS/XA, SRM, SNA/VTAM.
4 days/\$250 per day Sept 6 Hilton Head, SC Oct 11 Washington
(Sponsored by Duquesne Systems, Inc.)

2. MVS Overview.

2 days/\$500 Sept. 19 Dallas Oct. 6 Washington

3. MVS Systems Programming.

3 days/\$750 Sept. 21 Dallas Oct. 12 Washington

4. MVS Architecture (Structure & Logic).

3 days/\$750 Oct 3 Dallas Oct 24 Washington
Nov 14 Los Angeles

5. Exploiting the MVS/XA SRM Controls.

3 days/\$750 Sept 21 Dallas Oct 31 Washington

6. MVS Internals and Diagnostics.

4 days/\$1000 Dec 5 Dallas

7. VTAM Concepts & Facilities.

2 days/\$500 Oct 6 Dallas Oct 27 Washington

8. VTAM Internals.

4 days/\$1000 Dec 5 Washington

9. SAS® Basics Plus.

3 days/\$500 Sept 26 St. Louis Oct 17 Washington
Nov 14 Portland

10. Systems Awareness for Operators & Support Personnel.

3 days/\$750 Nov 7 Washington

"It's our instructors who make the difference. Once you attend one of our seminars, you will become a regular customer."

Bill Carico

PRESIDENT ACTS CORPORATION

Special Discount: Every National Systems Programmers Association (NaSPA) member is entitled to a special discount on each course. Call for details.
* SAS is a registered trademark of SAS Institute, Cary, N.C.

All MVS courses are current for MVS/XA version 2.2, All VTAM courses are current for ACF/VTAM version 3.1.1



(915) 388-3110

P.O. BOX 180
KINGSLAND, TEXAS
78639

The components of SAA are referred to as layers that will be implemented commonly across all operating environments. Currently they are SYSTEM 370, System 3X and the Personal Computer OS/2.

The components are:

- Common user interface
- Common programming interface
- Common communications support
- Common applications.

Common user interface

Common user interface will provide a common interface in three areas.

Syntax

There will be consistency over the location of information on the screen, for example, error messages always at the top, the sequence in which it is displayed and standardized action keys such as PF7, PF8 for scrolling and so on.

Semantics

The same command will produce the same action in whatever environment it is invoked.

Layout

The keyboards will have the same layouts and keys in all environments.

Common programming interface

Common programming interface will provide two features that will operate consistently in all environments.

Languages

- C, Cobol and FORTRAN for complex programming
- An application generator, Cross Systems Product (CSP), that, according to an IBM statement at SHARE, will probably be replaced
- A screen generator and manager, REXX (ISPF Dialogue Manager replacement).

Services

- Structured Query Language (SQL) for database access
- Query Management Facility (QMF) for report generation
- An EZ-VU type product for presentation services (EZ-VU is a PC product) and GDDM (graphics).

Common communication support

Common communication support has five components:

- Data streams (3270, Document Content Architecture, Intelligent Printer Data Stream)

- Application services (SNA Distribution Services, Document Interchange Architecture, SNA Network)
- Management architecture
- Session services (LU 6.2)
- Network Services (low level networking [PU2.1])
- Data link controls (X.25, Synchronous Data Link Control (SDLC), Token Ring).

Most of the basis for the communications component of SAA was taken from Systems Network Architecture (SNA).

Common applications

Common applications are applications conforming to the other SAA components that can be run in all the environments. A percentage of these common applications will be supplied by third-party vendors. It will also open up the possibility for IBM software to run under different operating systems.

Why SAA?

IBM has a major problem with transportation of products and information between environments. However, is that the only reason for SAA?

Isn't it interesting that the majority of products in the original announcement came from the small and medium environments? Where were IMS, CICS, PL/1? Is this another tack to grow the small shops from the DOS/VSE world or combat competition in the low end? Are large accounts not growing at the predicted 60 percent compound growth rate (SAA will increase CPU utilization) or is it just a rationalization of IBM's software development?

Small systems have generally offered a richer interface than large systems. This has resulted from the need of a consistent single user interface to the total environment, as opposed to the larger system needing machine cycles to process multiple simultaneous users.

SAA is badly needed to help the users with massive application backlogs and large maintenance requirements, but can it be achieved in today's diverse environments? With billions of dollars invested in the software, are users going to convert systems and would they want to?

Recent Inclusions

CICS and IMS/DC have recently been included in SAA. Would IBM's large customers let them be excluded? However, should they have been included if SAA is to succeed? Will CICS and IMS/DC be

able to take advantage of LU 6.2? Can they be recoded to support peer-to-peer communications (rather than host-to-terminal)?

When Will SAA Be Delivered?

Both of the latest inclusions in SAA (CICS and IMS/DC) have had their latest releases deferred from the original timeframe for general availability. Is this indicative of the time it takes to successfully modify complex software?

Considering the fifteen years taken to implement SNA, the basis of the communications support component of SAA, many people in the industry believe it will take until the next century before it is implemented. Think about the computing changes that will occur in that timeframe. Look back at the last fifteen years in computing. How much will SAA change in the same timeframe?

Conclusion

With the current development problems, diversity of systems (environments) and the myriad of products, SAA is the answer to the MIS director's prayers. However, changing the environment this dramatically will not be an easy task and will not happen overnight.

SAA is a statement of direction for long-range planners. At some point in the future, their applications will run in any environment encompassed by the architecture. As time passes, the architecture will be refined, but today's SAA gives the foundation on which to build. ☹



Phil Layland is a corporate consultant on DB/DC and applications development at National Advanced Systems that markets IBM-compatible mainframe systems. The company is headquartered in Santa Clara, CA (408) 970-1000.

Cross from page 25

Thirdly, the security benefit is derived because XMS allows CSA to be bypassed. The system hardware only provides MVS with 16 physical storage protect keys to work with, so MVS uses keys 0-7 to protect its own components. Keys 9-15 are reserved for use by V=R programs (programs that cannot endure paging and are loaded and executed in real storage in such a manner that the virtual addresses always map to equivalent real addresses). That leaves key 8 which is assigned to all other jobs running that do not fall into any of the other categories. Since each user has its own address space in MVS, this design of virtual storage protects each private area completely. However, whenever a key 8 user passes data through CSA, the CSA storage addresses are common to each address space and the only protection mechanism for users sharing CSA is the physical protect key. Therefore, all key 8 users are not protected from each other while their data is within CSA. XMS closes this security and integrity loophole by bypassing CSA altogether.

Cross Memory Terminology

XMS terms may appear at time to be intimidating. However, they represent the logical first building block in laying a foundation for further study on the topic. Take a moment now and complete the Cross Memory Services Terminology Quiz provided with this article. Once the correct answers are identified, review the terms a few times and give them a chance to sink in. Developing an understanding of the terms will help level-set for the next article that will begin fleshing out the details. ☉

ABOUT THE AUTHOR

Bill Carico is president of ACTS Corporation which specializes in providing advanced systems training for the MVS environment. Carico has developed and taught classes for ACTS since 1981. He was formerly employed by IBM where he worked as an instructor teaching MVS diagnostics and internals. ACTS offers both public and on-site classes for MVS, VTAM, DOS-to-MVS conversion and SAS. To receive a public class schedule and course descriptions contact ACTS Corporation, P.O. Box 180, Kingsland, TX 78639, (915) 388-3110.

Data from page 44

Encryption products are easily implemented for stand-alone use and are executed by a single CMS-like command. As with the CMS Copyfile command, with VM1's SafeGuard the user has the option of encrypting one or more files at a time according to filename, filetype and/or filemode, so that an entire minidisk can be encrypted in one command. The encryption process can also be used to back up files to tape, writing an encrypted version of a file (or files) directly to tape in one step. Decryption is handled with the same command and with the same options. Users also have a choice of single or double encryption with one or two keys in alphanumeric or hexadecimal notation using either the Data Encryption Standard algorithm or a proprietary algorithm.

Implementing most encryption products is also easy. Most can be installed in less than twenty minutes with or without the participation of a systems programmer. Access control packages are considerably more time-consuming to install. RACF, for example, takes three to four days to install depending upon the skill levels of the people involved and the amount of resources available.

Whereas typical access control packages require substantial planning and ongoing administration (to account for all of the permutations of user categories, function categories, resource categories, schedules and access points), the encryption process can be performed by individual users on their own files or by the systems administrator inserting the command into EXEC2 and REXX programs to automate the function.

Furthermore, there is no need for concern regarding the operational status of the program. While complex access control systems can become dysfunctional, encrypted files remain encrypted and the encryption function itself exists only as a module that can be easily restored.

If cost is an issue, and the cost-effectiveness of any software product usually is, DP managers should consider the fact that the commonly used access control packages sell in the \$25,000 to \$30,000 range. Encryption products give full data protection for between \$5,000 and \$10,000. So it is easily within the budgets of companies that are using access control, smaller companies that previously could not afford such protection and companies acquiring 9370s.

Conclusion

The VM professional pondering the need for securing data on his or her VM

system should consider the following situation:

- You release a tape containing sensitive data and send it through an overnight service only to be told the next day that it has been lost or delivered by accident to another company that may or may not be a direct competitor. Should you and your management be concerned? If the data on the tape is encrypted, your only concern would be the day's delay.
- An employee in your data center is fired and, bearing a grudge, steals a tape from the tape library. He holds it as ransom, demanding his job back. If the data on the tape is encrypted, it is harmless and the tape can be recreated.
- A government agency seizes the tapes in your data center for evidential material. If the data on the tape is encrypted, the company cannot be harmed.
- A private investigator working for a law firm or a competitor successfully bribes an employee to obtain either printouts or copies of tapes. If the data on your system is encrypted, these people will be foiled.
- There is a leak of sensitive information to the press or your competition and senior management points the finger at the data center. If your data is encrypted you can prove to them that the leak was not your fault.
- Security has been breached resulting in a multi-million dollar loss to the company. A shareholder lawsuit is filed against the company charging the board of directors with failure to take adequate protection of corporate assets. By encrypting the data, you could prevent the loss and the lawsuit.

In conclusion, the trend today is toward stand-alone encryption products. With prices coming down below \$10,000 they should become widely used in all data centers for both on-line and off-line (tape) data protection. VM users cannot prevent access to their data files, but they can and should prevent the abuse of that access. Encryption is the necessary course of action. ☉

ABOUT THE AUTHORS

Steve Halperin is a writer who specializes in IBM mainframe issues. Steven Rev has more than 10 years of VM experience at all levels and acts as a security consultant for large financial institutions.

Tower Systems' EPIC/VSE and EPIC/CMS Keep DASD Response Times Down

By John Kador

As DASD networks have grown in both size and complexity, the need to keep response times down has emerged as a critical issue of MIS management.

In meeting a data center's performance objectives, a key element is eliminating DASD bottlenecks before they grind processing to a halt. Response times for business-critical applications are at the mercy of an organization's DASD network. Controlling a shared DASD configuration only becomes tougher as the number of datasets and drives increases.

Managing DASD response times involves two distinct but interrelated processes. First, a data center must be able to detect and correct instances of DASD response time degradation. This process focuses on the technical causes and methods of solving a failure due, for example, to Rotational Position Sensing (RPS) delays. The second process assumes that technical failures are part-in-parcel with organizational deficiencies. It details the organizational procedures by which DASD configurations and dataset placements are determined. Solving the general problem of DASD response time management involves attention to both the technical and procedural aspects of the issue.

For DOS/VSE and VM/CMS shops, EPIC/VSE and EPIC/CMS from Tower Systems International offer comprehensive DASD resource management facilities. Except for the environments in which they operate, both EPIC systems offer

similar features. While most of EPIC's capabilities are common to the dozen or so other DASD resource management products available on the market, the underpinnings of EPIC are different.

For a DASD resource management tool to work, it must be able to directly measure DASD response times and other data needed to detect and correct aberrations. The underpinnings of most DASD tools are similar to IBM's approach: they hook into the operating system to trap SVCs and replace operating system code at the time of system initialization. This technique is certainly serviceable, but it consumes significant amounts of CPU overhead. Moreover, products that use this method often make themselves release-dependent on the operating system. That is never desirable because it complicates maintenance.

In designing EPIC, Tower Systems wanted to make it completely independent from hooks into the operating system internals. To that end, EPIC does not trap SVCs and does not issue privileged key zero commands. Installing it requires no changes to the operating system software, application programs or to Job Control Language. A centralized Dataset Name (DSN) catalog contains all the information for the access and control of both tape- and disk-resident datasets.

The product incorporates Generation Dataset (GDS) control functions which allow the system to maintain any number

of dataset versions on either tape or disk. As generations expire, they are automatically rotated into the available resource inventory and remaining active generations are aged by one cycle. It supports multimedia dataset categories, where specific datasets may reside on a combination of disk or tape device types. EPIC/VSE uses CICS in order to display resource status.

EPIC/VSE is a fully-integrated control system for managing tape and disk resources. An automated system for optimizing the control and security of datasets, it is designed to boost response time and to eliminate lost data files, computer reruns and unexpected job cancellations that can result from manual allocations. It is available for both VM and VSE shops.

EPIC/CMS improves control and management of CMS modules files and minidisks. The components of EPIC/CMS help VM users economize on disk and tape requirements for backups, improve security and reduce the time needed for data scheduling. Released in October, 1987, EPIC/CMS consists of four integrated components that can also stand alone. The first component is Backup that controls VM/CMS minidisk and CMS file backup and restore. The second one is Archive that provides a facility to control, compress and move seldom-used modules between a CMS user's active data area and system archive areas. Tape, the third

Maximize the Efficiency of Your Tape Library

Clean up with **FATS & FATAR**

FAST ANALYSIS OF TAPE SURFACES (FATS)[™]

FATS is a tape utility program which permits online tape certification, verification and erasure.

- FATS is designed to certify tapes where they will be used, on the actual tape drives.
- FATS allows an unlimited number of tapes to be processed on each tape drive in a single job.
- FATS certifies scratch tapes and verifies that data files can be read.
- FATS supports high-speed data security erasure of tape data.
- FATS provides a permanent record of each tape tested in the form of comprehensive detail and summary reports.

FAST ANALYSIS OF TAPE AND RECOVERY (FATAR)

FATAR is a tape utility program which analyzes and copies tapes, especially tapes with data checks or bad data.

FATAR is designed to:

- Read tapes, including multi-volume and/or multi-file tapes.
- Recover from I/O errors on tapes which cannot be read by normal operating system data management.
- Copy tapes, including multi-file and/or multi-volume tapes in one execution.
- Modify data on tape (while copying).
- Display and correct data check blocks on tape.

- Validate and correct the record format of files on tape (including variable spanned records like SMF).
- Print any or every data block on the tape — in EBCDIC, hexadecimal, or both.
- Convert ASCII data to EBCDIC, EBCDIC to ASCII.
- Copy tapes to a different label type or device type, such as ANSI-to-IBM labels, or 6250 BPI to 3480 cartridge.
- Re-block tape files while copying.

For A **FREE COPY** of the FATS
and FATAR REPORTS and a
FREE 90 DAY TRIAL
CALL (201) 890-7300



Available for VS1, MVS and MVS/XA

**INNOVATION
DATA PROCESSING**

INNOVATION PLAZA
275 Paterson Avenue
Little Falls, NJ 07424-1658

CIRCLE #97 on Reader Service Card

component, provides management and reporting facilities to control tape drive usage and tape volume mounts. Last, Console Control provides VM operator messages and reader, printer and punch queues.

Drive Time, Delay Time

There are two major components to DASD response time: drive time and delay time. Drive time represents the time it takes for a storage device to physically position the head at the desired track. Drive time can be further divided into two pieces: latency (time required to wait for the desired record to pass under the read/write head) and transmit time (the time required to transfer the data to or from the track). Managing these times is relatively easy because they depend only on the dataset placements on the drive.

Delay time is the more difficult component to manage. Transmitting data to or from the drive requires a path from the drive to the channel. If a path is not available at the instant needed, the connection must be delayed as the track rotates one complete revolution. Only then can another connect attempt be made. This is called Rotational Position Sensing (RPS) delay time.

Still another type of delay time is queue time. It occurs when an I/O is ready to be issued, but either no path to the drive is available or the drive is already executing an I/O. Moreover, in a shared DASD environment, delay time can be generated on a given dataset by I/Os from any of the CPUs connected to the drives and drive paths. Delay times are much more difficult to manage than drive times because delay times mostly reflect the overall DASD activity rather than the activity of a single dataset. Often it is impossible to pinpoint a single dataset whose activity is causing a delay.

Critical Issues Warrant Conversion

According to Olan Mills Studios in Chattanooga, TN, tape management is a critical issue that warranted the trauma of converting from one tape management system to another. Olan Mills, a nationwide photography studio, converted from a popular DOS tape management system in favor of EPIC/VSE.

"Service and support were the key reasons for the switch," says Pete Clark, systems programmer and data communications administrator for Olan Mills. The Olan Mills data center operates an IBM 3083 under DOS/VSE and an IBM 3033 VSE/VM and 88 DASD devices. The company manages more than 12,000 volumes of tape.

Performance differences (between the old system and EPIC) did not enter into Olan Mill's decision to convert. "For us, there were no performance issues. In terms of performance, we felt that both products were equal. We could not have justified the change on performance alone," Clark continues. EPIC is slightly faster and Olan Mills did gain back about 300K in core storage but, as Clark says, "that was not enough to make a change or to be an issue in the decision." Of course, the 300K available for other uses is appreciated.

Hughes' Dilemma

In 1985, the Connecting Devices Division of Hughes Aircraft Corporation in Irvine, CA faced a dilemma. It wanted to convert its Burroughs system to a 4381 under VM/HPO 4.2 and guest DOS/VSE systems. In doing so, it wanted immediate control of its DASD resources which included three strings of 3375 tape drives and a half string of 3380 disk drives.

"EPIC/VSE gave us a transparent method to keep on top of our datasets and

minimized our workload tremendously. "It has helped control the number of personnel in the I/O area, giving us more throughput and productivity without a lot of effort," says Steven Kannard, manager of technical support.

In strategic terms, Kannard believes better management of dataset resources resulted in decreases of CPU utilization of 40 percent. Those savings, when combined with productivity gains seen by many EPIC/VSE features, can cost-justify the product, he points out.

"Kannard continues, "During the conversion process, it was very nice to be able to treat unlabeled tapes off the IBM as though they were labeled for purposes of tracking datasets." The Burroughs system did not recognize unlabeled tapes very well. Consequently there were a lot of files getting passed back and forth between the machines. Using EPIC/VSE, we were easily able to track files leaving the IBM system and to determine when they were created."

The Connecting Devices Division also uses EPIC/CMS to manage DASD for its VM/CMS users. "It gives us the advantage of saving incredible amounts of time on the CMS side of the house," Kannard says, adding that it improved control and management of CMS modules files and minidisks by allowing all tapes created under VM to be controlled out of the DOS/VSE catalog. "Our ultimate goal is to have a single integrated catalog to control both VM/CMS and guest environments," he emphasizes.

Of EPIC/CMS, Kannard notes the partial backup and automatic recovery features. "Being able to do partial backups really made a difference for us," he adds. After the initial backup, the system backs up only files that have been updated since the previous backup. The archive facilities have been enhanced, as well. If users inadvertently erase one or more of the CMS datasets, they can now initiate the recovery procedure without involving the usually overworked operations staff, according to Kannard's tests.

For more information about EPIC/VSE or EPIC/CMS, contact Tower Systems International, 2220 Fairview Road, Costa Mesa, CA 92627. Telephone: (800) 854-7551, (714) 752-8263. ☉

ABOUT THE AUTHOR

John Kador is a free-lance writer and a frequent contributor to MAINFRAME JOURNAL.

CACHE IN with FASTIO V S E

ADD NEW LIFE TO YOUR CPU * IMPROVE CICS RESPONSE TIME


- * INCREASE SYSTEM THROUGHPUT BY 20% *
- * CUT BATCH RUN TIMES UP TO 50% *
- * SAVE UP TO 70% OF DISK I/O *

MAXIMIZE YOUR DOS/VSE SYSTEMS PERFORMANCE WITH FASTIO
CALL US NOW FOR YOUR FREE - NO OBLIGATION - 30 DAY TRIAL

1-800-752-FASTIO


GENERIC SYSTEMS INCORPORATED
2123 MANGO PLACE
JACKSONVILLE, FLORIDA 32207

(Prices and References Available Upon Request)



System 38, 4300, 9370 Users:

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



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

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

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

more storage capacity than an IBM 3422.

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

1-800-338-8ipl

(in MA 617-890-6620)

ipl systems, inc.

360 Second Avenue, Waltham, MA 02154

IBM is a registered trademark of International Business Machines Corporation

CIRCLE #96 on Reader Service Card

Forum from page 8

VSAM SORT WORK AREAS

I was informed by a number of readers that IBM's sort product supports only one extent of the type "SD" and nine extents of the type "DA." What this means is that IBM's sort does not support secondary allocation when using VSAM for the sort work area. However, all the other benefits of having your sort work area in VSAM still apply. At the New England Journal of Medicine we use CA-SORT from Computer Associates which supports multiple extents of the type "SD." I have been informed that Syncsort also supports multi "SD" type extents.

The second question I am most frequently asked is why I use explicit defines for the sort work clusters. I use explicit defines because of the reduced amount of I/Os and duration of the sort. When defining an explicit file with "no allocate" specified, the cluster is put in the catalog, but no space is allocated until a job opens it for output. This can be controlled by the "DISP=(.)" parameter on the dlbl card.

When you define an implicit file, VSAM loads every module known to man to do the define at the time the job opens that file for output and then performs the define.

In my article, "VSAM Tuning of System Work Files," in the November/December 1987 issue, I showed the results of benchmarks on an idle system using IBM's "PT" and job accounting information when using explicit

defines versus implicit defines. The results are significant and worth the effort to set up the clusters as explicit.

Finally, if you are using IBM's Sortmerg product, you should contact your IBM SE to request a PASR (product application and requirements) for the Sortmerg product to support multiple "SD" file type extents. This will allow you to take advantage of the secondary extent allocations when using VSAM/SAM, sort work area.

*Terry Spaulding
The New England Journal
of Medicine*

NON-UNIQUE AIX KEYS

I have an answer to the letter from Ronald D. Williams of Pansophic Systems, Inc. that appeared in the May/June 1988 issue about non-unique VSAM AIX keys and pseudo-conversational CICS transactions. My preferred technique is to avoid using the PATH and process the AIX and CLUSTER as separate datasets. Read the AIX for the desired key, save it in your COMMAREA and use program logic to extract the desired CLUSTER key. Then you can use that key to read the CLUSTER directly. I still define the PATH for documentation and for use by batch programs, but I ignore it in CICS. If you insist on using the PATH, you can save the last key used in the COMMAREA, do a series of READNEXTS

until you get a match on the last key used and resume processing with the next record. I wouldn't recommend this technique since the AIX can be read as easily as any other VSAM KSDS dataset. The fields in the AIX record are of known length and fixed (that is AIXkeyCLkey1CLkey2...CLkeyn) and are easy to process.

*David E. Taylor
Methodist Hospitals*

"X"37 ABENDS

Responding to Sue Tremblay's letter that appeared in the May/June 1988 issue regarding "X"37 ABENDS . . . two years ago, we also looked at STOP-X37 from Empact Software as a solution to serious production interruptions. Up-front analysis revealed that our users were allocating across multiple volumes, creating enormous fragmentation and performance problems. Products such as STOP-X37 were created to address problems caused by inadequacies in IBM storage management design and, as such, have usefulness. Our response was user education that resulted in immediate remediation of over half of our problems. Now that storage is cleaned up and users know how to estimate resource and appropriately allocate datasets, STOP-X37 may be a reasonable alternative to move us toward automated operations. Unfortunately, as long as IBM storage management is awkward, we

RENT IBM

4300 PERIPHERALS

Hamilton puts all the advantages of renting right at your fingertips. Waiting for IBM shipments? The solution is as easy as one call to us. Our rentals solve capital fund constraints and many other short term needs everyday. Hamilton is the answer!

CALL TOLL-FREE AND ASK FOR MR. JEFF BULL:

Eastern U.S. 1-800-387-7216 Canada 1-800-268-0317 Toronto (416) 252-6171

**HAMILTON**
COMPUTER SALES AND RENTALS

More ways to meet your needs.

New York, Houston, Dallas.

Montreal, Ottawa, Toronto, London, Edmonton, Calgary, Vancouver.

must educate, monitor and analytically correct allocation problems utilizing "after the ABEND" tools only if we understand what is causing the ABEND.

Suzanne Nott
Pacific Telecom, Inc.

APPRECIATES VARIETY

I just wanted to let you know how informative I find MAINFRAME JOURNAL. The great variety of topics discussed has broadened my base of knowledge in systems programming. In my area of concentration (VTAM, NCP, NetView), I rarely get involved with other IBM system products, but was always interested in understanding how they interrelate. Your magazine has allowed me to do this in a timely and interesting fashion with real-world examples and without the usual drudgery that accompanies a search through vendor manuals.

Ralph Coulter
PMA Group

NETVIEW/PC ARTICLE

I always enjoy reading about SNA network management; therefore, I was pleased to see the article entitled "SNA Network Management Using NetView and NetView/PC" in the May/June 1988 issue of MAINFRAME JOURNAL. However after reading it, I felt that a couple of points should be cleared up.

The IBM solution offered by NetView/PC is seen by most third-party vendors as an inadequate attempt at providing an interface to the NCCF network management platform. Third-party vendors with whom I have spoken feel that IBM could open up the interface and provide support for a number of different vectors in the NMVT RU in addition to the alerts that are currently supported in the NetView/PC interface. I share the concern that IBM holds the key to the network through network management and that the NMVT alert function is inadequate to meet most third-party network management needs. A few third-party vendors have written their own network management applications that take advantage of the NCCF macro level interface. These applications are fairly easy to write for an experienced IBM software engineer.

In addition, I noticed that when the authors were explaining the function of the CNM routing table, their solution to solve a CNM unsolicited RU routing problem would not accomplish the desired results. The article suggested that the NMVT NS RUs could be directed at the NEWS ACB by varying the BNJDSEV ACB inactive and by varying the NEWS ACB active. In ACF/VTAM, an APPL definition is an ACB. When an APPL is varied active, it is left in CONCT status. This status indicates that the APPL is available for an ACF/VTAM program to OPEN. Until the application opens the ACB, the ACF/VTAM APPL definition is not active and NS RUs will not be routed to it.

BNJDSEV is the NPDA task in NCCF.

When a CNM task is started in NCCF, the ACB is opened. When the NCCF task is shut down, the ACB is closed and the associated CNM ACF/VTAM APPL is returned to CONCT status. It is true that ACF/VTAM will not route NS RUs to two destinations, but there is no problem with having multiple destinations active at the same time. ACF/VTAM will route unsolicited network management RUs to the first "active" ACB entry found in the table for that RU.

In the example in the article, varying the BNJDSEV ACB inactive and the NEWS ACB active would have no effect on the NS RU

routing. Varying the BNJDSEV APPL inactive could cause problems later when you try to recover NPDA. Instead of varying the BNJDSEV APPL inactive, it would be better to use the NCCF STOP FORCE = BNJDSEV command to shut down the NPDA task. When the task is terminated, the ACB is closed by NCCF and the ACF/VTAM APPL returns to a CONCT state. Then, if the NEWS application is running and had opened the NEWS ACB, the NMVT NS RU would be routed to NEWS instead of BNJDSEV. At this point, the user could reassemble the table to recover the RECFMS RU.

Continued ►

XA-RELO / XR-RELO

CICS Performance Optimizer

As you may have already discovered, converting to XA does not necessarily mean your CICS performance and storage problems are over. Now it's time to let the powerful features of XA-RELO provide the solutions. (XR-RELO provides similar solutions for MVS/SP).

- Enjoy faster response times
- Improve internal throughput
- Eliminate all CICS storage compressions
- Eliminate virtual storage constraints
- Eliminate Short-on-Storage conditions
- Increase the Dynamic Storage Area (DSA)
- Eliminate all program fetches from the CICS load library during execution
- Reduce system I/O and WAIT time
- Allow all programs and mapsets to reside in the XA address space without any recompiles or modifications, including macro level programs
- Easy to install ... less than 30 minutes without any system modifications or program changes

— Call now for a free trial —

(800) 542-7760 • FAX (205) 854-0727



Quantum International Corporation

"Superior Solutions"

MHtran-1 VSE to MVS

The most advanced VSE to MVS migration software... anywhere!

For all the details or a FREE demo, call: "The Migration Experts"

MHT Services, Inc.

65 East Route 4, River Edge, NJ 07661
800-MHtran-1; In NJ: 201-342-1321

DOS/VSE and CICS/VS Frustration?

BIM gets it out of your system.

BIM presents a line of proven programs that maximize your system's capabilities, saving you time, labor and expense. These program products help get the most out of your system and people.

- BIM-EDIT** — the editor with more than 25 significant features that ICCF can't match.
- BIMSPPOOL** — Prints output in POWER/VSE spooling queue on local or remote 3270 terminal printers. (Received ICP Million Dollar Award 1982).
- BIMSPPOON** — On-Line to Batch Print Spooling. Prints data passed from CICS application programs into the POWER spooling queue.
- BIM-PDQ** — POWER Dynamic Queuing performance enhancement. Eliminates 85% of the I/O to heavily used POWER queue.
- BIM-ODIS** — Comprehensive problem analysis and display of operational CICS system. **DOS** and **OS**.
- BIMTEXT** — Word processing, document composition system. Create formatted documents from free-form input.
- BIMWINDOW** — Multiple terminal sessions concurrently at CRT under **DOS** or **OS** VTAM.
- BIMSWAP** — Switch local 3270 BTAM terminals between multiple CICS partitions without special hardware or additional ports.
- BIMCMPRS** — CICS 3270 data compression system. Reduces response time for remote terminals significantly. Available for **OS/VS1** and **MVS** also.
- BIMECHO** — Copies one CRT's output to another or printer for problem determination and demonstration.
- BIMP3270** — Comprehensive CRT screen image print facility. Copy to terminal printers or spool queue for system printer.
- BIMSERV** — On-line display of library directories and entries, VSAM Catalog entries, disk VTOC's, etc.
- BIMCNSOL** — Multiple/Remote System Console function for CICS. Display-only or full input/display versions available.
- BIMMONTR** — DOS/VSE System Status, Performance Measurement, and POWER Queue display.
- BIMSUBMT** — On-line Job Edit and Submission facility.

BIM programs are cost-efficient, some \$1000, highest \$5600. You can save even more with our group package offerings. Products are available on permanent, annual, or monthly licenses, and shipped on a 30-day free trial basis. Product documentation is available on request.

BIM also performs systems programming consulting, with consultants based in Minneapolis and Washington, D.C. Computer time services are also available on our 4331-2 system, on-site or remote.

BIM

B I MOYLE ASSOCIATES, INC.

5788 Lincoln Drive
Minneapolis, MN 55436

612-933-2885

Telex 297 893 (BIM UR)

Member Independent Computer Consultants Assn.

There is clearly no need to vary ACBs or ACF/VTAM APPLs active or inactive to control NS RU routing. In fact, that procedure would result in no network management data being collected and a potential recovery problem with NPDA.

Lloyd Hagemo, Jr.
BlueLine Software Inc.

VSE . . . Really!

I really enjoyed Pete Clark's "Really!" article because he is really correct with really good points about the real world of VSE/SP. It's a reality that any systems programmer, given adequate resources and training, who can't make VSE run really well, couldn't make MVS run either.

Let me add Conversion Reason #6: "To improve one's resumé and marketability." The demand for large-scale VSE systems programmers & VSE tech managers is small compared to MVS systems programmers.

We increased our workload and throughput with VSE/SP 2.1's bigger Power blocksize, multiple-page datasets (4), new Librarian, VAE, and 4K paging along with DASD tuning and 3380s.

VSE/SP 3.1's new Power will make it fly and reduce I/Os. VSE/SP 2 or SP 3 really runs fast . . . VSE 1.3.5 or SP 1! Really!

Arthur Stipo
Delaware Trust

IMBED Caution

The VSAM article by Frank Berezny (May/June 1988) was excellent. However, it should be mentioned that IMBED should not be used in a *cache controller* environment. A track of a replicated single SSI is "staged" to cache memory when many (if not all) SSIs could be staged instead. Also, IMBED causes secondary (non-SSI) devices to be "fragmented" with extent data when a VSAM file goes to physical extents. Since IMBED uses about 7% of the space it can actually be a cause of file extents in 3380 DASD.

Howard Glastetter
State of Washington

KUDOS

As the head of a large shop, I find the technical material in MAINFRAME JOURNAL superb. I grew up worrying about the dispatcher on OS/PCP release 2 when it was only doing one thing at a time anyway! I found the "new ready queue" comments very interesting; I hadn't realized ROTATE had disappeared in XA. The VSAM tuning article (by F. Berezny) was excellent and I have a JCL task force using it as a guideline.

Martin Thorpe
U.S. Administrators

Excellent magazine, how about increasing the size or frequency.

Jose L. Flores
Arden Mayfair, Inc.

Publisher's Note: Thanks for the kind remarks. We will be increasing the frequency to monthly starting in 1989.

Water from page 38

the fetching process begins all over again. Program compression can be reduced by taking some simple steps.

First, if your page size is set at 4096K you might try reducing it to 2048K. This will yield savings by reducing what I call page waste. For example, all non-resident programs are loaded on a page boundary. If you have a lot of programs that are less than 2048K in size, they will still consume an entire 4096K page if your page size is set to that. More than half the page will be "wasted." If you change the page size, you create more pages (twice as many) and you have more efficient use of available pages. By making this one change, we reduced program compressions from one every minute or so to one every six minutes.

In evaluating resident programs consider that programs specified resident through the Application Load Table (ALT) are not loaded on a page boundary and are in fact packed together. By examining the CICS out list and looking at the fetch-and-use rates of your load modules, you can determine whether or not the program should be resident. If it should be resident, specify so in the ALT.

These few suggestions are very basic in nature, but they will give you a starting point. Joining an organization like the Computer Measurement Group (CMG) will give you access to the best performance management and capacity planning minds around. Publications from CMG and other sources will help you gain experience as you work and perhaps save you the punishment of mistakes that have already been made somewhere else.

There will be frustrations in gaining support for starting a program especially where such a program has never existed. There will also be frustrations in analyzing and solving certain performance problems. Think about this: any rock, no matter how hard or how rough, when thrown into running water will eventually relent to the gentle, persistent and endless action of the water and become smooth. When these frustrations arise, do not quit. You must be like water on the rock. ☸

ABOUT THE AUTHOR

R. Douglas Swords has been in data processing for eight years and is currently the technical services manager for Pitt Memorial Hospital (Greenville, NC). His technical thrust for the past three years has been performance management and capacity planning.

SYSTEM PERFORMANCE

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

-MVS -VM -VSE -VTAM
-CICS -IMS -VSAM -COPICS

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

DOWLING ASSOCIATES, INC.

Information Systems Specialists

OHIO
(216)
656-1136

CONN.
(203)
255-6316

NEW YORK
(212)
912-1212

FLORIDA
(305)
782-9522

VSAM - EDIT

Realtime Dataset Editor and Management Facility

- Edit any VSAM dataset with CICS and/or TSO
- Automatically generate, centralize, maintain and submit all VSAM definitions
- Eliminate VSAM definitions from production JCL
- Perform real-time DELETE and DEFINE functions
- ADD, UPDATE, DELETE, EDIT, SEARCH, BROWSE, or DISPLAY any VSAM dataset
- Display data records with their COBOL field names
- Accurately calculate dataset space for any device
- Eliminate the conversion effort to ICF catalogs

— Call now for a free trial —

(800) 542-7760 • FAX (205) 854-0727



Quantum International Corporation
"Superior Solutions"

CICS ELECTRONIC MAIL



SYSM[®]

- Automate Internal Communications
- Leverage Existing Network Investment
- Designed for Non-Technical Users
- Full Function (In/Out Basket, Filing, Electronic Forms, Mail Lists, etc.)
- Well Over 400 Customers

Affordable Full Function System

- \$7,995 DOS • \$14,000 OS
- EasyLink Interface Available
- Calendaring/Scheduling Available
- PC Support Available
- DISOSS Interface Available



[®] H&W COMPUTER SYSTEMS, INC.
P.O. Box 15190 • Boise, Idaho 83715-0190
(208) 385-0336 • 1-800-338-6692
FAX: (208) 342-5219

DB2 from page 12

queries are predefined through application programs has required us to optimize as much as possible to achieve production-required response times. This has led to some interesting discoveries.

One of the virtues of the SQL language is that there are several ways in which a query may be phrased, all of which will return the same result. Theoretically, they should all perform equally well. However, with Release 1 we found that the insertion of parentheses in an "anded" condition reduced the response time for a single query from minutes to nearly instantaneous while obtaining the same result (this was subsequently corrected in Release 2). We have found many other similar cases which still persist. The message is that, while it is theoretically not necessary (and the requirement is decreasing over time), detailed knowledge of the DB2 optimizer and how it operates is important for applications designers to have.

When a query is executed, logically, a result table is created. Whether you wish to process only ten or a million rows of the result table, all rows which are candidates for the result table are determined (even if it is 5 million). DB2 addresses the formation of the logical result table in one of two ways: either by creating the entire result as a virtual table before presenting it to the user, or by retrieving each row of the result table from disk as the user requests them. The former is the best for larger queries where all of the result table is desired and the latter is better for smaller queries. Knowing this, the user may be able to fool DB2 into creating or not creating a result table in order to improve performance.

The construction of indexes is very important for many reasons. Locking is generally a greater problem through the index pages than through data pages since there will be far fewer index pages than data pages. Updating a single data row can lock many data pages through an effective lock on the index page which contains an entry for the updated data page, as well as entries for many other data pages. Also, only one index may be specified as the clustering index (the one which determines the physical order of the data on disk). However, through careful planning, several clustered indexes may be created.

DB2 is very good at using as much information in an index as is available in order to solve a query using only the index without having to go to the data page.

However, only one index may be used on each table in a given query. For example, the COURSES table will require an index on the Course# column to guarantee uniqueness (we assume that the DE-

PARTMENTS table has a unique index on the Dept# column). In addition, you may wish to add an index on the Dept# foreign key in the COURSES table for joins. See DB2 page 88

VSAM-LITE

Automatic Data Compression for VSE

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

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

CALL...

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


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

XA-VSAM / XR-VSAM

Improve VSAM performance as much as 90%

- Cuts hours off daily production time
- Significantly reduces I/O and WAIT time
- Analyzes and dynamically tunes the performance characteristics of all VSAM datasets
- Automatically provides optimum VSAM buffer management for maximum efficiency
- Eliminates dedicated resources now used to tune VSAM datasets
- Requires no program or system modifications
- Easy to install . . . less than 30 minutes

— Call now for a free trial —
(800) 542-7760 • FAX (205) 854-0727



Quantum International Corporation
"Superior Solutions"

Questions are answered by consultants and instructors from Davis, Thomas & Associates (Minneapolis, MN), the largest technical services firm in the upper midwest. Please address your technical questions to: The Tech Advisor, MAINFRAME JOURNAL, PO Box 38185, Dallas, TX 75238.

Q It is my understanding that VSAM gets the Control Area size from the primary or secondary allocation unit or 1 cylinder, whichever is smallest. What does VSAM do if the secondary allocation is zero?

We have several files that have performance problems and many of them have no secondary allocation amounts like *CYLINDERS (53 0)* and *CYLINDERS (45 0)*. We were wondering if the lack of a secondary allocation made VSAM pick a small Control Area size (like a track) and if that is causing some of our performance problems.

A That's a good question. You identified an exception to the primary/secondary/1-cylinder rule.

If a secondary allocation for a dataset is zero, VSAM ignores it when assigning the Control Area size. VSAM then selects the Control Area size by choosing only between the smaller of the primary allocation or 1 cylinder. (In an FBA environment the primary allocation is compared with 744 blocks.)

The two examples you gave would result in Control Area sizes of 1 cylinder which tend to offer the best performance for large files. You can verify the Control Area size on a LISTCAT report by looking at the field called "TRACKS/CA." It should equal the number of tracks per cylinder for your disk device. Sorry, but the answer to your performance problem lies elsewhere.

You may want to analyze why you don't have any secondary allocations on your files. You could have files that have been over-allocated to prevent them from running out of space. If not, you may be in danger of them filling up and causing problems for you later.

Q How difficult would it be to make VSE support a 31-bit as opposed to 24-bit addressing scheme? MVS seems prohibitively expensive for mid-range shops. The 16 MB (40 MB with VAE) seems to be the problem with VSE while overhead of CPU, disk space requirements and so on limit VM.

A To make 31-bit addressing available for DOS/VSE may be impossible in my opinion. The problem is in the usage of registers in the software. Many technical packages and system software use the high-order byte of the registers to pass information and condition codes. It may require a rewrite of VSE, POWER, CICS, VTAM and other vendor packages. I don't think IBM is going to introduce DOS/VSE/XA either.

Pete Clark, systems programmer and data communication administrator for Olan Mills, Chattanooga, TN has enhanced VSE and extended its capability well beyond Release 3.1. If you would like a copy of Clark's code, circle number 390 on the Reader Service Card.

As for running DOS/VSE under VM, I don't recommend running it in VAE mode. DOS/VSE/VAE under VM loses many of the VM handshaking capabilities along with increased paging caused by double paging and also a loss of the real memory dedicated to a DOS/VSE/VAE system to help it run effectively.

I recommend that DOS/VSE run under VM with a *MODE=VM* supervisor and a directory entry of 16 MB minimum and maximum size specified. VM will then do all of the paging. Multiple VSE guests can then be controlled by the type of work being performed by each.

Q What are the performance considerations of: *SET QDROP OFF*, *VIRTUAL=REAL*, *SET (CCW) NOTRANS ON*, dedicated disk and full pack minidisks for VSE/SP?

A When a Virtual Machine is being dropped from a scheduling queue, the machine's segment and page tables are scanned to determine which pages should be flushed. Setting *QDROP OFF* prevents the machine's pages from being flushed and as a result improves performance for the virtual machine. Although the

VSE/SP machine may see a performance improvement, other virtual machines may experience adverse effects. An alternative to setting *QDROP OFF* is to lock strategic pages (that is, page 0, CICS tables and control blocks) for the VSE/SP machine. This is an area of tuning in which trial and error is needed to determine what will work best for your situation.

A VSE/SP machine running in VAE mode must be defined *VIRTUAL=REAL* to VM. When running *VIRTUAL=REAL*, the VSE/SP machine's pages are mapped 1 to 1 in real storage (except for page 0). VM supports only one guest machine running *VIRTUAL=REAL* mode. A VSE/SP Supervisor that does not have the VM handshaking support is used and, as a result, does its own paging. One performance benefit is that CCW translation can be omitted by issuing a *SET NOTRANS ON* command for the VSE/SP machine. A major disadvantage is dedicating a portion of real storage to one virtual machine that could leave other virtual machines contending for the remainder of real storage. My suggestion is to run multiple VSE/SP machines in *VIRTUAL=VIRTUAL* mode unless there is a major advantage to running your VSE/SP in VAE mode.

I do not know of any major performance issues concerning full pack minidisks vs. dedicating DASD volumes. One functional benefit of full pack minidisks is the ability to share DASD between multiple VSE/SP systems, although I do not recommend sharing DASD unless absolutely necessary.

Q Regarding MVS, on the EXEC statement, what is the effect of using the *DPRTY=(,)* parameter? How does the SRM react to such an alteration in practical terms, for batch processing and for a single step?

A The effect of using the *DPRTY* parameter on the EXEC statement depends on several parameters in 'SYS1.PARMLIB' which define the Automatic Priority Group (APG) range. The APG range is the range of dispatching priorities that is under control of the System Resource Manager (SRM). The dispatching priority determines the position on the dispatch queue on which a task is placed when it is ready to execute.

Two cases determine the effect of the *DPRTY* parameter:

Case 1: the *DPRTY* parameter values set the dispatching priority if the first value specified *does not* fall within the APG range. From the *DPRTY=(value1,value2)* parameter the priority is set to a 2-digit hexadecimal number where the first digit is obtained from decimal 'value1' and the second digit is obtained from decimal 'value2'.

Example — *DPRTY=(8,15)*:

Decimal '8' = Hexadecimal '8'.

Decimal '15' = Hexadecimal 'F'.

Therefore, the dispatching priority = Hexadecimal '8F'.

Case 2: the *DPRTY* values are ignored if the first value specified falls within the APG range. In this case, or if *DPRTY* is not coded, the dispatching priority is obtained from the specifications in PARMLIB member IEAIPSSx, also known as the Installation Performance Specification (IPS).

The parameters that determine how *DPRTY* is used are:

- APG=nn* in member IEASYSxx sets the Automatic Priority Group value which is used as the first digit in the dispatching priority
- APG=mm* in member IEAIPSSx sets the second digit in the dispatching priority for a specific performance group
- APGRNG* in member IEAIPSSx sets the range of dispatching priorities that are controlled by SRM
- DP* in member IEAIPSSx sets the dispatching priority for a specific performance group.

See Advisor page 92

How to Buy Time for Your IBM 308X.



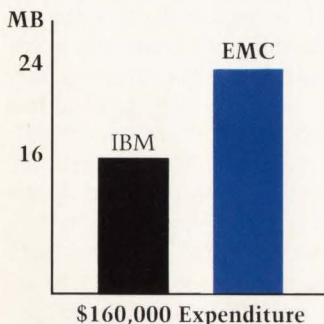
At 30% Off.

1. Invest in EMC.

With the uncertainty surrounding IBM's next move in the mainframe area, adding a few more years to your system's productive life will pay dividends for you down the road. So why buy a new system, when EMC's main storage upgrade gives you the performance you need to make your current system your future workhorse. With EMC you save 30% over comparable IBM upgrades, which means buying time now for your 308X has never been a better decision. The simple "plug and play" design of EMC's upgrade makes improving the performance of your system as easy as adding sand to an hourglass.

2. Maximize Your Savings.

EMC uses state-of-the-art technology and production methods to offer our 308X upgrade at 30% lower cost than IBM. A 16MB upgrade from IBM is priced at \$160,000. For that same price you could buy 24MB of EMC main storage. That's an additional 8MB of main storage or a savings of \$50,000. Either way you save with EMC. EMC also features coterminous leases and trade-up credits so the return on your investment continues after your purchase.



3. Improve Performance.

By increasing your 308X's main storage capacity, you will improve system response time and increase your users' productivity.

EMC's 308X upgrades are 100% hardware and software compatible with IBM 3081, 3083, and 3084 CPU's. Upgrades are provided in 16MB increments and run all IBM diagnostics. Use of EMC upgrades will have no effect on your IBM maintenance.

4. Add Reliability.

To guarantee the most uptime possible for your 308X system, our upgrades use pre-tested logic components. Then they are subjected to 100-hours of test and burn-in procedures, which include qualification in one of our own 308X computers prior to shipment.

EMC is the leading independent manufacturer of main storage upgrades and other system enhancement products for mainframe and mid-range computers. All EMC products are supported by our worldwide network of sales and service offices.

5. Call EMC.

If you're interested in buying time for your 308X system, call EMC today at **800-222-EMC2** for your free information kit. Our 308X upgrade is the "midlife kicker" that saves you money and improves your 308X's performance.

For more information, call today:
1-800-222-EMC2 ext. 2290
(In MA, call 617-435-1000)

EMC²

The System Enhancement Company.

IBM is a registered trademark of International Business Machines Corp.
Copyright 1987 EMC Corporation

CIRCLE #127 on Reader Service Card

DB2 from page 85

Say that we also have made the Course# index the clustered index because it is frequently qualified in queries and also provides a good physical distribution of data. Given the query in Figure 5, one of the two indexes may be used, but not both. Thus the data page will have to be reference for every row in the index. Use of the clustered index would be preferred. However, if we add the Course# column to the Dept# index, now the entire join portion of the query could be

solved using only the index. However, it is not the clustered index, so data retrieval would be in effect in random sequence.

One final change to the indexes might be to make the clustering index be the combination of Course# and Dept# (in that order) and another unique index on Course#. Of course the clustering index would have to be non-unique, but the result is that both indexes will be clustered. Now the query will be solved in much the same way as above, except that the data

is in the same order as the index.

As you can see, knowledge of your queries and the index structures is quite important (note that all of the above changes to the indexes could be made and utilized by applications without any impact or required change to their code).

FIGURE 5

```
SELECT *
FROM   COURSES, DEPARTMENTS
WHERE  COURSES.Dept# = DEPARTMENTS.Dept#
```

A simple join

One other problem situation may frequently arise when two tables are compared (whether as a part of a join or a simple restriction of a join) on other than primary and foreign key columns. When columns of dissimilar data types (though both are numeric or both are alphanumeric), the query takes a long time to solve. Were domains properly supported, this problem would be alleviated. (Note: it is assumed that the problem will not exist for primary-foreign key comparisons because the table designer has ensured that they have the same datatypes.)

FIGURE 6

```
SELECT *
FROM   COURSES, DEPARTMENTS
WHERE  COURSES.Dept# = DEPARTMENTS.Dept#
AND    DEPARTMENTS.Dept# = 111

SELECT *
FROM   COURSES, DEPARTMENTS
WHERE  COURSES.Dept# = DEPARTMENTS.Dept#
AND    COURSES.Dept# = 111

SELECT *
FROM   COURSES, DEPARTMENTS
WHERE  COURSES.Dept# = DEPARTMENTS.Dept#
AND    DEPARTMENTS.Dept# = 111
AND    COURSES.Dept# = 111
```

Join with restriction moving from worst to best

In addition, the way information is provided to DB2 is important. Say the restriction is added that we wish only to see information for Dept 111 in the query above. There are three ways to code this request (Figure 6), all of which are sufficient. The last is the best, despite the fact that the information provided by the two restrictions seems completely redundant. However, it allows DB2 to apply the restriction to both tables prior to attempting the join. In future releases, we are confident that providing such redundant information will not be necessary. The second form of the query which is presented is also better than the first (can you determine why?).

CICS and VSAM SOFTWARE

\$495 - \$695

FREE TRIAL - OVER 5,000 SOLD

SUBMIT CEMT COMMANDS FROM BATCH

CICS/CEMT from BATCH - OPEN or CLOSE CICS files from a batch job stream. Initiate CICS tasks, issue any CEMT SET command. START/STOP DL1 data bases (DOS), ALLOCATE & UNALLOCATE files (MVS). Works with up to 99 CICS systems on multiple CPUs. Over 1300 users. *Purchase for \$495 (DOS) or \$695 (MVS).

SEND MESSAGES TO CICS TERMINALS

CICS/MESSAGE - Send messages to CICS terminals without affecting user's transactions. CICS/MESSAGE saves the screen buffer, Tran-id and COMMAREA and displays the message. The user's screen and transaction are fully restored. Messages may be sent to one terminal, a set of terminals, all terminals, or the operator console. CICS/MESSAGE may be linked to from other CICS programs. *Purchase for \$495 for both DOS and MVS.

SHOW WHAT IS ON ANOTHER CICS TERMINAL

SHOW AND TELL - Show what is currently on another user's screen. Great for helping users without having to go to their office. TELL will echo your screen to other screens or printers. TELL can be used to print screens and for demonstrations and training sessions (local and remote). The SPY transaction (included) allows you to monitor a user's terminal. A 3270 print key function is also included. \$895 for purchase.

ELIMINATE IDCAMS LISTCATS

LISTCAT PLUS - Recapture wasted disk space, save time and paper by using LISTCAT PLUS. Lists VSAM entries (including ICF VSAM) in one-line-per dataset format. Datasets needing attention can be flagged. Lists non-VSAM entries in ICF catalogs, too. Over 1300 users. Purchase for \$495 (DOS) or \$695 (MVS).

*Effective September 1, 1988 the prices for the following products will increase. CICS/CEMT (DOS) - \$695, (MVS) - \$895. CICS/MESSAGE (DOS and MVS) - \$695.

Mackinney Systems

2674-A South Glenstone, Suite 112 Springfield, Missouri 65804 (417) 882-8012

Over 25 other products to choose from for CICS and BATCH. All products are available for 30-day free trial. Annual lease prices are also available.

Another knotty problem is that of the introduction of redundancy into a table. All of our tables are the products of a rigorous logical design process which produces normalized tables. Normalization is a discipline which guarantees that the data in each table is fully and solely dependent on the primary key for that table. This avoids many potential update anomalies and eliminates redundant data. This is the preferred state for a table design to be in. However, due to limitations in the DBMS' ability to handle more than a few tables in a single query with adequate performance, it may be necessary to carry some columns in more than one table. This avoids the need for a join in some cases. The trade-off is that this redundant data must now be maintained. We have found that a judicious introduction of redundancy can help with performance. However, we have been quite satisfied in most cases with a fully normalized set of tables.

Finally, one feature of DB2 is that it allows for queries to be "stored" as "views." Views may then be referenced in other queries as though they are tables, removing unwanted complexity from the concerns of the users. In addition, views can be used to create restrictions or projections for the purpose of security. This is accomplished by only allowing access to the view and not to the base table on which the view is based. In many cases this has proven to be quite satisfactory. However, not all views are updatable. In fact, only views referencing a single table in their primary query may be updated (in DB2). In order to obtain the desired security through views with these, a subquery is added which may join to the other restricting table. Unfortunately, this

may substantially impact the performance of the query since the subquery must be evaluated for all possible rows, not just candidate rows. Again, if few rows of a large table are to be updated, the impact could be substantial.

As you can see, performance is directly affected by the knowledge that the designers have of SQL and DB2. We have found that unacceptable response times are more often than not caused by poorly coded queries. Overall, with careful at-

tention, DB2 has proven itself to be a true production DBMS. ☉

ABOUT THE AUTHOR

David Allen Pollock is senior technical consultant for the data management department at Ranier Bank Services, a Security Pacific Company, in Seattle, WA. He was graduated with a bachelor of science and information systems degree from the California Polytechnic University, Pomona, CA.

HOW TO LOSE WAIT FOR SUMMER:

Get new VM1 QuickCopy and Copyfile up to 15 times faster.

Summer's here, so what better time to reduce excess system wait? With new VM1 QuickCopy™, you can copy, pack and reformat CMS files up to **fifteen times faster** than CMS Copyfile.

What's more, QuickCopy speeds up your entire system, not just copy operations. Because all of your applications will benefit from QuickCopy's CPU time savings.

QuickCopy does everything Copyfile does, plus some handy extras. And because it uses the identical syntax and return codes, users don't have to learn any new commands and existing programs don't have to be modified. Just plug it in and go...fast!

To help you start your wait reduction plan, VM1 is offering QuickCopy at a special introductory price of only \$3995.

So get your VM system in shape for summer. Give QuickCopy a 30-day trial at no obligation and see how much wait you lose. And how much productivity you gain.

Call 1-800-321-4VM1
(in CA, 213-641-2600)

5250 W. Century Blvd.
Los Angeles, CA 90045



*Remember,
information is
available for all
products advertised.
Just fill out the
READER SERVICE
CARD elsewhere
in this magazine.*

Candle Fills Vacuum in VM Performance Reporting

Candle Corporation's new EPILOG 1000 for VM provides historical reports for VM users. It also provides reports analyzing DASD and channel resource usage systemwide, as well as CPU, storage and paging usage both systemwide and by workload. EPILOG 1000 supports VM/SP, VM/HPO and VM/XA. This follows IBM's recent announcement that the system reporting tool for VM, VMMAP, will not support VM/XA.

For more information
CIRCLE #200 on Reader Service Card

ADR's MINDOVER MVS Improves MVS & MVS/XA Performance

Using artificial intelligence techniques, ADR's MINDOVER MVS improves MVS and MVS/XA performance by performing detailed analysis of the MVS system, identifying performance problems, finding their cause and suggesting solutions to improve performance. It enables performance specialists to implement solutions instead of spending countless hours pouring over volumes of MVS performance data.

For more information
CIRCLE #201 on Reader Service Card

Advantages of CICS to VM Development Now

MAGEC Software's new VM MAGEC simulates CICS terminal handling, provides activity logging, a chronological history of transactions and audit stamps in VSAM records. VM MAGEC provides the advantages of CICS to VM development and the ability to transparently move to VM/CICS later.

For more information
CIRCLE #202 on Reader Service Card

4GL/Compiler Expanded to Support SQL/DS

Phoenix Software has extended the capabilities of its fourth-generation language/compiler, CYGNET, to support IBM's SQL/DS database language in DOS/VSE (SP) environments running on-line under CICS, VTAM and VSE (SP) batch. CYGNET/SQL is reportedly the first application development tool of its kind to deliver 4GL productivity improvements plus assembler-level run-time performance.

For more information
CIRCLE #205 on Reader Service Card

New Utility Automatically Converts DFHSM Datasets

Innovation Data Processing has announced the availability of a conversion utility that will automatically convert datasets migrated under DFHSM to FDRABR archived datasets. Installations seeking the outstanding performance and reliability of the FDRABR DASD Management System will be able to move data now under DFHSM control to FDRABR control easily with this new conversion utility.

For more information
CIRCLE #206 on Reader Service Card

Software DASD Caching Available for VM/XA Users

Cache Magic/XA, just announced by SDI, is said to improve the performance and increase the processing capacity of a user's system, specifically in a VM/XA environment. With Cache Magic/XA, installed users can experience dramatic reductions in batch run times, accelerated on-line response times and an increase in overall throughput due to the system's ability to process more transactions in the same period of time.

For more information
CIRCLE #207 on Reader Service Card

New Amdahl Front-end Communications Processors

The new 4745 Series FECPS from Amdahl allow migration from ACF/NCP-4 to ACF/NCP-5 software and reportedly give users up to 40 percent more throughput than competitive products. Because they run industry-standard software, the new 4745 Series processors can be intermixed in networks with Amdahl's earlier 4725 and 4705 models as well as IBM's 3745, 3725, 3720 and 3705 communications controllers.

For more information
CIRCLE #208 on Reader Service Card

DTA Releases Print Package for DOS/VSE and VM

DTA/Print Version 3.0 has been released by Davis, Thomas & Associates as a comprehensive report distribution and viewing system that allows any report in any IBM Power or VM Spool Queue to be routed to any attached CICS printer or CRT. Security features allow Power and VM file protection, plus system and printer control.

For more information
CIRCLE #203 on Reader Service Card

First Totally Interactive Tool for CICS Dumps

SymDump, the first productivity tool that reportedly allows application programmers to easily resolve any CICS transaction dump, was recently introduced by On-Line Software International. SymDump lets programmers examine the desired information in a symbolic format; thereby, eliminating the need to translate memory offsets and hexadecimal values.

For more information
CIRCLE #209 on Reader Service Card

CA-SPOOLMAN Enhances VM System Management

Compiler Associates International's CA-SPOOLMAN is a recently announced spool management facility that provides a wide variety of automated functions that are not available in the standard VM environment. CA-SPOOLMAN provides real-time monitoring and extensive viewing capabilities of VM spool activity and provides users with maximum control while operating in the VM environment.

For more information
CIRCLE #210 on Reader Service Card

First VM Tool for ASCII/EBCDIC Data Exchange

VM1 has introduced CrossLink, a new software tool that enables VM users to exchange data with non-IBM computers such as DEC, H-P, Data General, etc. Not available through any VM feature offered by IBM, CrossLink is believed to be the first product of its kind to solve the problem of ASCII/EBCDIC incompatibility.

For more information
CIRCLE #212 on Reader Service Card

New Performance Monitors for CICS/MVS & VM/XA(SP)

Goal Systems International recently announced EXPLORE/CICS for MVS and EXPLORE/VM for IBM's new VM/XA(SP). EXPLORE/CICS for MVS helps data center management economically monitor CICS performance and also monitors performance across operating systems allowing performance analysts to produce one report with information from all company data centers. EXPLORE/VM for VM/XA(SP) allows users of the current VM operating system to continue to perform monitoring and capacity planning activities. When they convert to VM/

XA(SP), the migration will be transparent from a performance point of view.

For more information
CIRCLE #213 on Reader Service Card

DB2 & IMS Users Get New Menu-driven Query System

DYL-IQ EXPRESS from Sterling Software's Dylakor Division is an interactive query and immediate response reporting system that combines a menu-driven facility, natural language syntax and other approaches to accessing information in DB2, IMS and VSAM data structures. DYL-IQ EXPRESS provides immediate access to corporate data structures simultaneously and is available in three versions: CICS, IMS/DC or TSO telecommunications monitors.

For more information
CIRCLE #214 on Reader Service Card

MICS VSE/POWER Analysis New From Morino

MICS VSE/POWER Analysis, from Morino Associates, provides a complete picture of batch workload activity for organizations running a combination of MVS and VSE operating systems (or with VSE remote). These users will also gain a unified, enterprise-wide view of batch activity on their systems. By interfacing VSE/POWER Analysis with Morino's MICS Accounting and Chargeback Product, I/S organizations can produce a single, consolidated invoice reflecting the mixed services of MVS and VSE.

For more information
CIRCLE #215 on Reader Service Card

CICS Programmers Get FAST-EDI

D & E Software has announced FAST-EDI, a CICS utility that provides easy access to VSAM, ISAM and DL/I files on-line. It was designed as a CICS development and debugging aid and for technical support personnel for the interrogation and correction of corrupted files.

For more information
CIRCLE #216 on Reader Service Card

Integrated Collection of Tools for DB2 Users

DB Observer, from North 40 Software, Inc. addresses problems of application performance, management of complex databases, database design and efficient use of DASD space and CPU time. DB Observer is said to reduce the cost

Additional detailed information on all products described in Product Update is available by circling the number designated at the end of each product description on the Reader Service Card.

of designing, coding, maintaining and administering DB2 applications and to help users create smaller and faster programs.

For more information
CIRCLE #204 on Reader Service Card

New DOS/VSE (SP) VSAM Disk Compression Product

BIM-PACK is a new program from BI Moyle Associates, Inc. that compresses DOS/VSE(SP) VSAM disk files automatically (transparent to programs that access the files). Use of BIM-PACK is said to result in the reduction of disk space occupied by the files compressed by 35 to 70 percent.

For more information
CIRCLE #211 on Reader Service Card

Enhancement to RD/SHARE Aids VM Users

RD/SHARE Release 4.0, from RD Labs, extends the comprehensive CMS library management facilities to include a menu-driven full-screen user interface, an EXPAND facility that dynamically resolves references to include files and extensions of version control capabilities to accommodate variable length records.

For more information
CIRCLE #217 on Reader Service Card

DB2 Users Can Now Execute Batch Without TSO

DATABASE ATTACH is a new product from Softbase Systems that allows users of DB2 to execute batch programs natively without TSO. The system uses the Call Attach Facility of DB2 to communicate with the database; thereby, eliminating the need for TSO to execute batch programs.

For more information
CIRCLE #218 on Reader Service Card

Instant Feedback for DB2 Developers

Relational Architects, Inc. has released RLX/DS that provides a direct SQL interface to DB2 from the interpretive TSO CLIST language. RLX/DS eliminates the preprocess, compile, linkedit and bind steps that typically delay DB2 application development.

For more information
CIRCLE #219 on Reader Service Card

New Cross-Reference & Documentation Package

Computer Concepts has announced its new cross-reference and documentation software package for MVS users — the QXR-

19. QXR-19 has the capability to access and quickly research libraries containing source, object modules, procedures, JCL, load modules and user-defined records such as selected information concerning CICS transactions.

For more information
CIRCLE #220 on Reader Service Card

New Utility for Validating JCL Syntax & Standards

JCLAUDIT from Consumer Systems is said to offer the most automated and cost-effective way to provide maximum assurance that OS/MVS JCL is free of errors before the "production" run. It should substantially relieve job re-runs due to JCL errors and help ensure compliance with corporate JCL naming conventions.

For more information
CIRCLE #221 on Reader Service Card

PANAPT Automates Production Turnover Management

Pansophic's PANAPT fully automates the program production turnover process required to identify, approve and move applications into production libraries. It also manages the migration of application entities from TEST to QA to PRODUCTION and provides a comprehensive on-line history of the entire promotion process.

For more information
CIRCLE #222 on Reader Service Card

ASF Automates Scheduling

The Automated Scheduling Facility (ASF) from Chaney Systems Support was introduced to solve the problem of job scheduling in an MVS/JES2 environment. ASF includes: ISPF/PDF split screen support, dynamic creation and maintenance of calendars, unlimited calendars, end-user scheduling capability, separate scheduling and execution controls, centrally controlled security and many other features.

For more information
CIRCLE #223 on Reader Service Card

Cybermation's DJC Bridges IBM's CJS Gap

After IBM dropped support for Chained Job Schedulers (CJS), Cybermation announced that its Dependent Job Control (DJC) allows continued use of CJS control cards without the need for one single change. DJC is said to also offer CJS shops many extra benefits. It can be used to control complex

relationships between jobs, providing JES2 installations with the same dependent job control capabilities of JES3.

For more information
CIRCLE #224 on Reader Service Card

VSE Users to Benefit from SAS Version 5

SAS Institute has announced the SAS System under VSE, an integrated software system for data management analysis and presentation. Version 5 includes integrated modules for data entry, retrieval and management; report writing and graphics; statistical analysis; business planning, forecasting, and decision support; operations research and project management; and an interface to IBM's DL/I databases.

For more information
CIRCLE #225 on Reader Service Card

ROSCOE Users Get Testing and Debugging Tool

Application Development Systems, Inc., in cooperation with ADR, has just announced availability of XPEDITER/ROS, a full-screen interactive testing and debugging tool for COBOL programmers who use the ROSCOE on-line development system from ADR (Applied Data Research). XPEDITER/ROS lets the programmer control execution of his/her program, step through code, display and modify data, intercept abends and watch execution using a full-screen source level display.

For more information
CIRCLE #226 on Reader Service Card

CL/SUPERSESSION Now Uses Less CPU Resource

Using the latest VTAM high-performance options, Candle Corporation's CL/SUPERSESSION is reportedly now as much as 17 percent more CPU efficient than the previous version of this MVS multi-session manager. Other enhancements in this version of CL/SUPERSESSION center around the Dialog Manager that allows control of screen presentation and processing logic. Now users can write more flexible and powerful dialogs for certain functions.

For more information
CIRCLE #227 on Reader Service Card

DASCOMP Revolutionizes Archival Backup

Information Management Software has announced DASCOMP, a software program that reportedly

compresses data for archival backup by as much as 80 percent with no additional hardware required. By compressing data up to an 8-to-1 ratio prior to transmission, DASCOMP is said to increase data transfer throughput between computers from 100 to 200 percent in far less time as well.

For more information
CIRCLE #228 on Reader Service Card

Westinghouse Introduces Integrated Network Solutions (WINS)

Westinghouse Management Systems Software has integrated its popular Network Control Interface (NCI) and Multiple Session Manager (MULTSESS) and expanded its network access and control product portfolio to produce a powerful, integrated network solutions package capable of running across all IBM operating systems. NCI, the core of the WINS, provides quick, easy access to applications through a powerful single point entry system. MULTSESS enables users to simultaneously access multiple applications from a single terminal, saving time and valuable hardware resources.

For more information
CIRCLE #229 on Reader Service Card

Online-SyncSort Now Available

Syncsort Inc. has announced Online-Syncsort, a new facility for using SyncSort OS interactively. Online-SyncSort is ISPF-based and generates the control statements necessary for SyncSort sort, merge or copy applications. The automatic generation of control statements eliminates the possibility of Syntax errors.

For more information
CIRCLE #230 on Reader Service Card

New GEN-A-PIC Generates Hardware Installation Diagram

Candle Corporation just revealed GEN-A-PIC for MVS, a new product that generates a hard-copy diagram of MVS/XA and MVS 370 system hardware configurations. It simplifies the process of adding, deleting or moving system devices from one channel to another by visually representing the relationship between control units, I/O devices (terminals, DASD, printers, PCs and tape drives), and the channels of a configuration.

For more information
CIRCLE #231 on Reader Service Card

Advisor from page 86

The two main options on how to use these parameters are to code or not to code APGRNG and DP:

Option 1:
Coding APGRNG and DP override APG in

IEASYSxx and is mutually exclusive with APG in IEAIPSxx.

Option 2:
If APGRNG and DP are not coded then the APG range is set by the APG parameters in IEASYSxx and IEAIPSxx (or defaulted).

Regardless of the option chosen, an APG range is either defined or defaulted.

Examples

1) APGRNG = (6-15) is specified in the IPS and a batch job has DPRTY = (15, 15) specified on the EXEC statement.

In this example, the dispatching priority would be obtained from the IPS specification for batch jobs. The DPRTY parameter would be ignored. (See Case 2 above).

2) APGRNG = (9-15) is specified in the IPS and a batch job has DPRTY = (8, 15) specified on the EXEC statement.

In this example, since the first value in the DPRTY parameter (8) is outside the APGRNG, the dispatching priority would be '8F'. (See Case 1 above).

3) APGRNG = (0-15) is specified in the IPS and DPRTY = (12, 15) is specified on an EXEC statement.

In this example, all DPRTY specifications would be ignored and the dispatching priority would be obtained from the IPS. This would be the case for all batch jobs regardless of whether or not DPRTY was specified.

4) APGRNG not specified in IPS. APG = 7 is specified in IEASYSxx. APG = 6 is specified in IEAIPSxx. DPRTY = (7, 15) is specified on an EXEC statement.

In this example, since 7 is the default APG range, the second value in DPRTY is ignored and the dispatching priority is set to '76'. The '7' is obtained from the APG parameter in IEASYSxx and the '6' is obtained from the APG parameter in IEAIPSxx. (See Case 2).

5) Same as 4) except that DPRTY = (12, 15) is specified on an EXEC statement.

In this example, since 12 is outside the APG range, the dispatching priority would be 'CF'. (See Case 1).

To ensure centralized control of dispatching priorities, I'd recommend the following:

- 1) Code APGRNG = (0-15) in IEAIPSxx so that all DPRTY parameters in JCL are ignored and the dispatching priorities are obtained via the IPS.
- 2) Set up multiple performance groups for batch jobs in IPS to accommodate the different batch priorities (that is, hot, production, test, long running, etc.).
- 3) Develop job naming standards, a batch initiator structure and a corresponding ICS (Installation Control Specification — member IEAICSxx in 'SYS1.PARMLIB') to control how jobs are assigned to a performance group and dispatching priority.
- 4) Implement routines to enforce the standards. They could be JES2, SMF, and/or TSO Submit exits.

(This answer was researched and written by Gene Heikkinen, systems consultant.)

THE

CICS

Series
by **Data 21**

ONLINE HELP

CICS.Help™

The original Help windowing package for CICS, which includes Application Windows. Two products for the price of one.

DATA COMPRESSION

Turbo Response™

50% to 85% terminal compression. Includes feature to view other terminals for Help Desk, Demos, etc.

ONLINE SPOOLING

Remote Print Facility™

Display/Print reports from system spoolers to 32XX devices. Five years of proven performance.

Data 21

CREATING A BETTER ENVIRONMENT

DATA 21 3838 Carson Street, Suite 121, Torrance, CA 90503 (213)543-5599

Q: What do 71 of the Fortune 100 companies have in common?

A: A cost effective data entry method utilizing ODE II's integrated on-line/PC system.

ODE II is state-of-the-art data entry software that replaces any key-to-disk, key-to-tape, or card system at a fraction of the cost. It was designed for data entry on an IBM Mainframe and IBM PC or plug compatible. A few of the benefits ODE II offers are:

- **Flexibility** to design applications on-line on the PC.
- **Compatibility** with any micro to mainframe link or emulation program.
- **Usability** through features such as built-in EDITS, subsecond response time and menu-driven approach.

ODE II is reasonably priced, provides flexible lease plans, low maintenance and offers a 90-day *unconditional* warranty after purchase. For additional information or to arrange a *no obligation* 30-day trial, call 1-800-633-4748.


ODE
ON-LINE DATA ENTRY

1112-7th Avenue
Monroe, WI 53566-9982
1-800-633-4748

IST International
Software
Technology, Inc.

MAINFRAME JOURNAL
An Official IBM System 370, 370/158, 370/158-1, 370/158-2, 370/158-3, 370/158-4, 370/158-5, 370/158-6, 370/158-7, 370/158-8, 370/158-9, 370/158-10, 370/158-11, 370/158-12, 370/158-13, 370/158-14, 370/158-15, 370/158-16, 370/158-17, 370/158-18, 370/158-19, 370/158-20, 370/158-21, 370/158-22, 370/158-23, 370/158-24, 370/158-25, 370/158-26, 370/158-27, 370/158-28, 370/158-29, 370/158-30, 370/158-31, 370/158-32, 370/158-33, 370/158-34, 370/158-35, 370/158-36, 370/158-37, 370/158-38, 370/158-39, 370/158-40, 370/158-41, 370/158-42, 370/158-43, 370/158-44, 370/158-45, 370/158-46, 370/158-47, 370/158-48, 370/158-49, 370/158-50, 370/158-51, 370/158-52, 370/158-53, 370/158-54, 370/158-55, 370/158-56, 370/158-57, 370/158-58, 370/158-59, 370/158-60, 370/158-61, 370/158-62, 370/158-63, 370/158-64, 370/158-65, 370/158-66, 370/158-67, 370/158-68, 370/158-69, 370/158-70, 370/158-71, 370/158-72, 370/158-73, 370/158-74, 370/158-75, 370/158-76, 370/158-77, 370/158-78, 370/158-79, 370/158-80, 370/158-81, 370/158-82, 370/158-83, 370/158-84, 370/158-85, 370/158-86, 370/158-87, 370/158-88, 370/158-89, 370/158-90, 370/158-91, 370/158-92, 370/158-93, 370/158-94, 370/158-95, 370/158-96, 370/158-97, 370/158-98, 370/158-99, 370/158-100

SHIFT YOUR PERFORMANCE INTO HIGH GEAR WITH CANDLE EDUCATION



Every day you rejoin the race to keep your system up and response time down. It's a challenge just to stay on top of all the latest developments.

You're on the fast track. So you have to team up with the best. Candle Education. Because Candle is the leader in performance management and data center solutions.

Our instructors are experts on the latest techniques for managing your MVS, CICS, IMS, or VM system. They also know all about three of the hottest new industry topics: networking, data base, and operations.

In fact, we offer over fifty classes designed to give you all the skills you need to stay ahead in this fast-moving race. That's why we can help make your job easier... and the rewards higher.

So call us at (213) 442-4180 for additional information, a course catalog and schedule, or to register. But hurry - classes fill up fast.

When the checkered flag comes down, make sure you're in the lead. Team up with the ultimate in high performance. Candle Education.

WASHINGTON D.C.

COURSE	DATES
MVS Fundamentals for Operations Personnel - MVS250	Jul 12-13
OMEGAMON for MVS for Operations Personnel - MVS255	Jul 14
Running MVS Under VM - Performance Considerations - MVS403/VM330	Jul 12-13
MVS I/O Subsystem Design and Modeling - MVS413	Jul 14-15
MVS for Managers - MVS700	Jul 19-20
IMS for Operators - IMS250	Oct 18-19
Operation of Candle IMS Products - IMS205	Oct 20
IMS Fundamentals - IMS200	Oct 24-27
Operation of Candle IMS Products - IMS205	Oct 28
Tuning an MVS System - MVS300	Oct 25-28
CICS Performance and Tuning for MVS - CICS300	Oct 25-28
High Performance CICS System Design for MVS - CICS400	Oct 25-28
VM/HPD Performance and Tuning - VM310	Oct 25-27
VM System Monitoring with Candle Products - VM305	Oct 28
ACF/VTAM Network Performance and Tuning - NET300	Oct 25-27
DB2 Performance and Tuning for MVS - DB340	Oct 25-28
High Performance VSAM System Design - DB400	Nov 1-4

BOSTON

COURSE	DATES
Operation of Candle MVS Products - MVS205	Jul 12-13
MVS System Monitoring with Candle Products - MVS305	Jul 14-15
Introduction to MVS Performance - MVS200	Aug 9-12
MVS Fundamentals for Operations Personnel - MVS250	Aug 15-16
Tuning an MVS System - MVS300	Aug 23-26

DALLAS

COURSE	DATES
Tuning an MVS System - MVS300	Jul 26-29
CICS Performance and Tuning for MVS - CICS300	Jul 26-29
IMS Fundamentals - IMS200	Jul 26-29
VM Performance and Tuning - VM300	Jul 26-28
SNA Fundamentals - NET200	Jul 26-28

NEW YORK

COURSE	DATES
VM Performance and Tuning - VM300	Jul 19-21
VM/SP System Monitoring with Candle Products - VM305	Jul 22
High Performance CICS Systems Design for MVS - CICS400	Aug 2-5
VSAM Performance and Tuning for MVS - DB300	Aug 9-12
SNA Fundamentals - NET300	Sep 20-22
Fundamentals of Performance and Internals for MVS - CICS200	Sep 27-30
Introduction to MVS Performance - MVS200	Oct 25-28
IPS Design/Performance Workshop - MVS415	Nov 8-11
ACF/VTAM Architecture and Programming - NET320	Dec 6-8

CHICAGO

COURSE	DATES
Fundamentals of Performance and Internals for MVS - CICS200	Jul 12-15
Operation of Candle CICS Products - CICS205	Jul 19-20
CICS System Management with Candle Products - CICS305	Jul 21-22
Fundamentals of Data Bases and Access Methods for MVS - DB200	Sep 6-9
VM System Fundamentals - VM200	Sep 13-15
Candle VM Product Operation - Fundamental - VM205	Sep 16
Data Base and Access Methods for Managers - DB700	Sep 13-16
Introduction to MVS Performance - MVS200	Sep 20-23
Tuning an MVS System - MVS300	Sep 27-30
IMS/DC Architecture and System Services - IMS300	Nov 1-3
CICS/VSAM Performance and Tuning - CICS520	Nov 8-9
CICS/VTAM Interface - CICS580	Nov 10-11
CICS/VS MRO/ISC Design Implementation and Performance Tuning - CICS500	Nov 15-18
CICS Fundamentals for Operations Personnel - CICS250	Nov 29-30
OMEGAMON for CICS for Operations Personnel - CICS255	Dec 1

LOS ANGELES

COURSE	DATES
ACF/VTAM Network Performance and Tuning - NET300	Jul 5-7
IMS/DC Performance Tuning - IMS400	Jul 11-14
IMS/DC Performance Management with Candle Products - IMS405	Jul 15
CICS/DB2 Interface - CICS560	Jul 14-15
CICS/VSAM Performance and Tuning - CICS520	Jul 19-20
CICS/VTAM Interface - CICS580	Jul 21-22
Operation of Candle MVS Products - MVS205	Aug 2-3
MVS System Monitoring with Candle Products - MVS305	Aug 4-5
CICS for Managers - CICS700	Aug 9-12
SNA Fundamentals - NET200	Aug 30-Sep 1
Fundamentals of Performance and Internals for VSE - CICS210	Sep 20-23
CICS Performance and Tuning for MVS - CICS300	Sep 27-30
Advanced Techniques for Predicting MVS Performance - MVS416	Sep 27-30
MVS Fundamentals for Operations Personnel - MVS250	Nov 8-9
OMEGAMON for MVS for Operations Personnel - MVS255	Nov 10
ACF/VTAM Architecture and Programming - NET320	Nov 8-10
VM Performance Fundamentals for Operations Personnel - VM250	Nov 15-16
OMEGAMON for VM for Operations Personnel - VM255	Nov 17
IMS/DB Performance and Tuning for MVS - DB320	Nov 29-Dec 2
High Performance CICS/VS Application Design and Implementation - CICS620	Dec 6-9
High Performance IMS/DB System Design - DB420	Dec 6-9
High Performance DB2 System Design - DB440	Dec 13-16

SAN FRANCISCO

COURSE	DATES
IPS Design/Performance Workshop - MVS415	Aug 2-5
CICS Performance and Tuning for MVS - CICS300	Aug 2-5
IMS/DC Performance Tuning - IMS400	Aug 2-5
VM Performance and Tuning - VM300	Aug 2-4
ACF/VTAM Network Performance and Tuning - NET300	Aug 2-4

Candle[®]



Candle Education, 1999 Bundy Drive, Los Angeles, CA 90025

Copyright © 1988 Candle Corporation. All Rights Reserved.

CIRCLE #131 on Reader Service Card

BMC Software, Inc.

All The Right Stuff in Leading-Edge Technology

Software happens.

This is a truism of the industry according to Ted Van Duyn, BMC Software's senior vice president of research and development. "Anyone or any company can develop software. The trick, then, is to make sure that it's the right stuff," he states.

Since its origins in 1980, BMC Software, Inc., has at times, only by the virtue of good timing, had the ability to bring the right stuff to the market place for the right price and at the right time. "Of course, there's much to be said about BMC's overall business philosophy when you look at its success with mainframe systems software and enhancements, but it generally comes back to the fact that BMC's management in marketing and research and development have had the insight to make listening to our customers a number one priority," Van Duyn says.

This exchange of information between BMC Software and BMC customers has developed from a tradition of continuous dialogue. "Customer feedback is critical not only to the software we develop, but also to the way it is marketed to professionals within the industry," says Max Watson, senior vice president of sales and marketing. "That is why we attend all the IBM User Group meetings and other conferences such as CMG. That's why BMC hosts several focus groups and management briefings throughout the year and that's the reason for our BMC Users Group (BUG) each year. Only by listening to customer requirements and responding to these needs will BMC maintain its reputation for product integrity and superior technology."

Van Duyn agrees, "We don't try to be flashy, we're not the playmakers of the industry. What BMC products do for the production system is much like the job of blocking and tackling in a football game. We deal with the technical aspects of computer operation. We're at the heart of data processing production and crucial to the efficiency of the operating environment."

BMC's beginnings were based on this same philosophy — delivery of superior products to enhance existing computer operating environments without fanfare. In 1979, BMC founder John Moores introduced his only software product, 3270 OPTIMIZER/IMS, first in and around Houston, and later, on a national basis. By the end of 1981, the tiny mainframe software company had had its first million-dollar year and as the product line has expanded, sales have skyrocketed.

Today, BMC Software products span five distinct product groups:

- Data Communications Series — reduces 3270 data streams to be transmitted through telecommunications networks, thereby saving resources and transmission requirements. This group includes BMC's OPTIMIZER for the IMS/DC, CICS and VM operating environments and SUPER-OPTIMIZER products for the IMS/DC, CICS and TSO operating environments.
- DASD Data Compression Products — reduce space required to store data on direct access disk storage devices for IMS, FAST PATH, VSAM and DB2.
- Database Utility Products — reduce resources and downtime during database reorganization for DOS/VSE and IMS/DB environments.
- DB2 Products — provide performance improvement for DB2. Products include: DATA PACKER/DB2 for compression of DB2 tables and DB2 ALTER that allows a user complete control of the change process for DB2 data structures.
- IMS/DC Enhancement Products — expand the capabilities of certain IMS/DC features allowing users to operate more efficiently within the IMS/DC environment.

All BMC products are designed to meet three specific criteria:

- To generate immediate and direct cost savings
- To operate transparently within a data processing system
- Are easily installable and supportable.

"BMC Software," says Watson, "has grown into a unique corporation that is immensely strong in:

- Responsiveness to requirements of users and potential users
- Investment in research and development
- Expansion to international markets
- Product support.

"However, the successful evolution of this corporation is no accident. BMC's philosophy requires a superior product backed by individuals who have superior abilities in the areas of service and support, as well as authors who are continually upgrading those products and our sales team, all of whom share the common goal of providing timely response to customer needs and requirements," Watson contends.

It has also been management's strategy

to hire the best in every field, provide an environment conducive to creativity and productivity and to make it possible for talented individuals to meet their own goals as well as the objectives of the corporation.


BMC's corporate headquarters, covering more than 100,000 square feet in Sugar Land, TX southwest of Houston, gives credence to this goal — spacious facilities, a smokeless environment and a casual working atmosphere with the open-door management approach. "We hire knowledgeable, experienced individuals and give them a place to do what they do best," emphasizes Rick Hosley, BMC's president and chief executive officer. "We've set some high standards and everyone here works to meet them, from mailroom to management."

Watson points out that because of BMC's uncompromising standards, it has become the premier supplier of performance-oriented systems utilities for IBM mainframe environments. "We try to make it easy for the customer to do business with us," he says. "We offer a 30-day-plus free trial, make support available 24-hours-a-day, 7-days-a-week and we don't require trial agreements or automatic billing. What our experience tells us, however, is that when a potential customer installs a trial of a BMC product, he immediately realizes significant benefits that he wants to continue after the trial ends. These trials usually result in sales."

BMC's marketing effort is supported by a series of Technical Reviews or one-to-one information sessions. "Our products require little, if any, training. What we try to provide is update information on the new technology that is available for more efficiency in a particular operating environment. This keeps our customers informed," Van Duyn points out.

Throughout its eight-year history, BMC Software has enjoyed mutually beneficial relationships with its customers. Hosley attributes this to the corporation's emphasis on dialogue. "Some companies don't have the time or resources to listen to their customers. At BMC, we're committed to our customers and to providing software that pushes ahead of what exists. We believe that many of our products go beyond state-of-the-art, if you will, simply because we listened to the requirements and concerns of our customers and then found a better way to address these areas," Hosley says. ●


Advertisers say the darndest things!



“**M**ainframe
Journal offers the
largest and most
targeted audience of
systems software
buyers and evaluators
for the DOS/VSE,
MVS, and VM
operating systems.”

Denny Yost, vice president of sales

BlueLine
SOFTWARE INC



For advertising information, call Denise Thomas (214) 343-3717

If...

VSAM performance and integrity is an issue.

Then...

Goal Systems International DCM-VSAM

The best supported VSAM product line in the industry.

/* FAVER—VSAM file management system with comprehensive reorganization and backup/restore capability. Available for MVS and VSE.

VSAMAID—Tune VSAM automatically while eliminating inefficiencies—a cluster modelling, performance and capacity planning tool.

Available for MVS and VSE.

MASTERCAT—An online and batch VSAM catalog information product for MVS and VSE.

COMPRESSOR/MVS—Provides compression of VSAM datasets that can triple disk capacity. */

To Find Quality Call 800-848-4640.



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

CIRCLE #79 on Reader Service Card