MAINFRAME JOURNAL

For Users of IBM System/370 Architecture & Compatible Systems

September / October 1988

New Coupon Marketing Concept Utilizes 9370 & VM

In this issue...

SNA's LU 6.2

3090 PR/SM Feature

CICS Storage Constraints

Electronic Report Distribution

VM Diagnostic Techniques

DASD Key Compression

and much more...

The VM Software Seminar gets you back on track.

Information strategy is supposed to work with corporate strategy, not against it. Yet in many organizations, strategic tug-of-wars make everyone a loser—while derailing company growth plans.

Fortunately, it doesn't have to be this way. And at the free VM Software Seminar, you can learn first-hand how to make information management more responsive to corporate needs—and less demanding on MIS personnel.

Systems management in the network era.

This unique half-day event offers a wealth of information on systems management tools and tactics that meet vital corporate objectives while simplifying administration in today's network-oriented VM environments.

You'll see the critical functionality that has made VMCENTER II the world's most widely acclaimed VM systems management package—a comprehensive resource for managing growing user populations amid escalating security and cost pressures.

You'll also get a glimpse into the future of network management in the form of a powerful new tool that promises to simplify control of your IBM SNA networks.

It's information you won't want to miss.

And it even comes with a free lunch. So don't delay. Make your reservation today. And come to the seminar that will help you keep your organization on the fast track—without getting derailed.

Seminar agenda

8:30 a.m.	Registration and Coffee
9:00 a.m.	Seminar Begins
12:30 p.m.	Complimentary Lunch

Reserve your place today, call:

VM SOFTWARE INC.

(800)562-7100
(703)264-8413

Seminar dates and locations

Boston, MA October 21	**Indianapolis, IN** October 5	**Raleigh, NC** November 16
Calgary, AB October 18	**Kansas City, MO** November 15	**San Francisco, CA** November 3
Chicago, IL October 6	**Miami, FL** November 17	**San Jose, CA** October 13
Columbus, OH October 25	**Minneapolis, MN** October 31	**Seattle, WA** November 1
Dallas, TX October 11	**New Brunswick, NJ** November 10	**Springfield, IL** October 4
Denver, CO October 12	**New York, NY** October 20	**St. Louis, MO** October 27
Detroit, MI October 26	**Orange County, CA** November 2	**Toronto, ON** October 19
Houston, TX October 28	**Philadelphia, PA** November 9	**Washington, DC** November 7
	Pittsburgh, PA October 7	

Note: attendance by a software vendor or its representative is prohibited without prior written approval from VM Software, Inc.

I want to stay on the strategic fast track. Please enroll me in the VM Software Seminar at_____ on_____.

Name_____

Title_____

Company_____

Address_____

City_____

State_____ Zip_____

Phone (_____)_____ Operating System_____

CPU-Manufacturer/Model_____

VM Software, Inc., 1800 Alexander Bell Dr., Reston, VA 22091

5-MFJ-8809

CIRCLE #20 on Reader Service Card ▲

Still screaming about why CICS failed?

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

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

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

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

LANDMARK™

eyewitness™

SOLVE THE MYSTERY

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

CIRCLE #29 on Reader Service Card ▲

CONTENTS

September/ October 1988

Volume III Number 4

DEPARTMENTS

COVER:

Cover design by Don Bryant for Adroit Designs Inc.

Photography by Robert Germany.

83 *Documentation*

FEATURES

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.

"Our network moves enormous loads quickly. So does our CICS."

"At CSX, we have 61 production CICS regions and seven million transactions per day. With The Monitor For CICS,™ we see problems long before our users do."

Rail transportation. Container shipping. Gas pipelines. Resorts. CSX is a $13 billion giant. With over 21,000 miles of rail, 6,000 miles of natural gas pipeline, and 5,000 miles of fiber optics, CSX needs real-time status to service its customers.

So at their Jacksonville, Florida, and Baltimore, Maryland, facilities, CSX uses CICS to track status and inventory—and relies on The Monitor For CICS to manage CICS performance. "The Supertrace feature lets us look inside an application and gauge its effects on system performance," says **Jason Butler, Manager, Technical Services.** "We can trace application logic and evaluate resource consumption right down to the event level."

"We've built a unique monitoring system that is PC-based and set up so that Monitor commands are automatically executed to identify poor response times. This lets me spend more time with features such as the storage display. Now when a problem arises in CICS, I can alter storage or delete ICE/AID chains rather than shutting down and cold starting the system."

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

LANDMARK™

THE MONITOR® FOR CICS

MINDS YOUR BUSINESS

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

CIRCLE #16 on Reader Service Card ▲

PUBLISHER'S COMMENTS

Robert H. Thomas

Our staff is extremely excited to present you with this issue of MAINFRAME JOURNAL. It is our biggest issue yet and it is loaded with a variety of excellent topics written by a top-notch group of writers. Enjoy!

Reader Forum

We've always prided our magazine on having a very intelligent readership and the letters in this month's Reader Forum starting on page 8 prove it.

Contiguous Articles

Recently we received several letters requesting that articles run contiguously, rather than being "continued" throughout the magazine. We took a shot at it this time.

Vendor Profile

A regular department in each issue of MJ is Vendor Profile. One reader indicated in a letter to us that he would like clarification of the purpose and intent of this department.

The vendors/companies that market their products by advertisements in our magazine may be known primarily by their logos and ad copy. Vendor Profile was implemented so that in each issue a different vendor could make a brief "presentation" introducing the company and its products to you in more depth.

PR/SM Causes Incompatibility

Just as we went to press with this issue, John Baker of Policy Management Systems called to inform us that IBM's PR/SM support for VSE/SP (APAR #DY36770) causes an incompatibility with Pete Clark's modification to VSE.

John and Pete conferred and reported that the incompatibility caused by PR/SM will affect few, if any. Even so, Pete has an update for VSE/SP 3.1.2 to resolve separation of VSIZE and VIO maximum allowable virtual space. Note that small additional displacement changes have occurred, so correct those also.

Add to fix #DY99001 the following code:

```
ALTER 0438 -
   0000A000F4F04000:00020000F1F2F800
/*    X'40' 4 0        X'128' 1 2 8                    */
/* REFERENCE NOT CONTAINED ON ANY CURRENT FICHE        */
/* APAR DY36770 REQUIRED BEFORE THIS CODE EXIST        */
/* MAX VALUE FOR VIO AREA SIZE, NOT REQUIRED FOR MORE ADDR SPACES*/
```

Speaking of Pete Clark, it's always a pleasure to have his comments on VSE and we have just that starting on page 21.

VM Diagnostics

Ed Sterling and Gabe Goldberg each present an article that should make life for VMers a little easier. Ed discusses CCWTRACE, the software tool that provides tracking information needed to diagnose difficult hardware and timing problems (page 102). In his article, Gabe offers a study in "Advanced VM Diagnostic Techniques" (page 106). These two "sleuths" can crack the hardest cases.

Robert H. Thomas

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

Publisher and Editor-in-Chief
Robert H. Thomas

Associate Publisher
Martha Thomas

Editor
Carol M. Hoag

Copy Editors
Judy Beller
Pat Warner

Art Director
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 USA and Canada. One-year foreign subscriptions are $96.

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

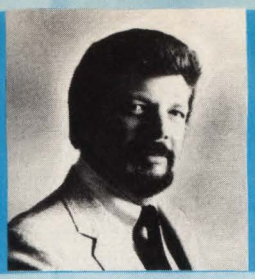

"It's time for a change."

Jerry W. Bennett • President • Bennett Software, Inc.

Bennett SOFTWARE

HOUSTON, TEXAS

PHONE:
1-800-562-8722 US
1-800-448-3336 CANADA
1-713-451-8832 FAX

August 15th, 1988.

To : All Clients and Prospective Clients.

From: J. W. Bennett
Subj: Maintenance charges, all Bennett products.

First, let me thank each of our current customers and soon-to-be customers for making Bennett Software such a startling success. In the 12 months beginning May 1st, 1987, our JOBTRAC control system captured an amazing 10% of new business for MVS job scheduling software in the U.S.. Funded only by a dedication to service and customer satisfaction, we surpassed the $1 million mark within our first year and are looking toward a $4 million second year.

We didn't acquire this success, we developed it. We develop every product we sell. We owe our success to the customers that chose us, not those who chose one of our acquired competitors.

It's time to break from tradition once again.

The MIS executives that I speak with every day are in agreement on one basic issue: The rising cost of yearly vendor software maintenance is becoming a genuine concern.

In the past, most software cost justifications focused on base prices and largely ignored the yearly maintenance costs. Vendors know this and exploit the issue. Some vendors impart 20% fees, or more. The more unscrupulous vendors will set base license prices at double what the product's are worth, then cut a discount deal off the list price by 50% to 75%. This assures them of a fat maintenance check every year, based on the nondiscounted list price.

Honest vendors require maintenance to pay for enhancements and developments that are due the customer, and keep the customer's investment abreast with the latest technologies. In theory, "maintenance" should not be used as a high profit "wrench" to use against customers. Recent maintenance increases, across the market, are causing major concerns in most data centers, especially when many of these products haven't been enhanced in years.

In an effort to force moderation from our competitors and give our clients relief, Bennett Software is announcing a reduction in yearly maintenance fees from 15% to **12%**, retroactive to January, 1988. All customers paying in excess of 12% during 1988, will be reimbursed. This maintenance level will be frozen for 18 months, or until March 1st 1990. Base initial license fees, for all current and newly developed products, will be fixed, for the same period. We will continue our practice of providing site licenses, rather than CPU licenses, throughout.

It's time for a change.

Sincerely,

J.W. Bennett

HOME OFFICE:
P.O. BOX 96694
HOUSTON, TX 77213
(713) 451-0191

ATLANTA:
24 HIGHWAY 213
COVINGTON, GA 30209
(404) 786-1686

WASHINGTON, DC:
8150 LEESBURG PIKE, #1252
VIENNA, VA 22180
(703) 448-6808

CIRCLE #61 on Reader Service Card ▲

READER FORUM

SOFTWARE MAINTENANCE — NO PROBLEM

This letter is concerning Jan Snyders' article, "Software Maintenance, Undermanaged & Understaffed," (July/August 1988). All the displayed points of view, as well as the recommended solutions for the so-called software maintenance problem are excellent; however, I really expected someone to represent the other side, the bright side, of the coin.

First of all, why do people call it a problem? Software maintenance is no different than any other maintenance task. When a construction engineer finishes a skyscraper, a maintenance crew is hired before people can actually use the building. This is not because people expect problems, but because it is a normal procedure in order to keep the building in the best operable condition at all times.

Another example is the continuous maintenance to your wardrobe. You may do or ask someone to do alterations, you buy new items or get rid of some old ones. You are doing this and other similar tasks continuously and without calling it a problem. Other examples of maintenance are your car, house, appliances and health. Maintenance is an ongoing process for the purpose of keeping the item continuously fit for the purpose it was created. This is necessary in overcoming the depreciation of the item as well as to help it adapt to any changes in the environment. In a couple of words, a good planner must allocate part of his budget for maintenance purposes.

If we think of software maintenance in that way, we could easily accept it as a normal and simple fact that follows the creation of any new software. More important, we could try to make it easier by planning, allocating money, managing and properly following up its activities. The fact that software maintenance is now affecting our lives more than any other maintenance job is simply because it may affect the decision-making process that leads to either a successful or failing business. This should make us alert to the same simple process of conducting effective maintenance which leads, with other factors, to a successful business. This process is planning, allocating money, good managing and following up the maintenance process.

The difference between types of maintenance is the availability of the level of expertise required for an effective and correct performance of the required maintenance from the first trial. Software maintenance may be one of the most difficult maintenances because of the obscurity of working through program code and logic that are sensitive to any changes because of their inter-relationships; nevertheless, it is a vital process for today's volatile business environment.

From my past experiences in this field, I have summarized a few software maintenance ideas. First, assigning specific persons to the maintenance job is a big mistake because it kills their creativity and they will not be as good in maintaining any newly developed applications. The staff performing the maintenance job must also share the application development process in order to maintain it or other related systems.

Second, maintenance tasks should be a part of each software programmer's or program analyst's daily activities. This part may range from 20 to 80 percent of his time and this percentage should not be held constant for long periods of time.

Third, maintenance of different applications must be rotated between staff members under the supervision of a more experienced person. This can help build the necessary experience required for effectively solving business problems.

Fourth, organizations should be careful in selecting analyzing tools; otherwise, these tools can be a burden on the maintenance staff. Fifth, organizations should establish and pursue certain standards for program coding. This procedure should start as early as the detailed design of a project. It will lead to easy and correct maintenance in the future.

Sixth, the cost of any new project should include, besides the one-time implementation cost, a yearly running cost which should be a certain percentage of the total project cost devoted to post implementation maintenance. This will keep the project fitting efficiently with the dynamic business environment.

Seventh, the 4GLs were mainly designed to help the end user and the non-EDP professionals get their required output from the computer. It will not replace the traditionally coded applications because, while it can easily handle the quick-and-dirty applications (as one of the viewers explained it), it cannot handle complex business and industrial applications.

Eighth, I would expect an expert system to exist in the market shortly that can write technical program specifications after scanning the program source. Such a product would be helpful for a person making changes to that program, because it can easily direct him to where he can effectively code the changes. Once the modifications are accepted, new program specifications can be documented using such product.

Shwikar Hadary
Gordon Jewelry Corp.
Houston, TX

VSAM TUNING

After reading Frank Bereznay's article on VSAM (May/June 1988), I was astounded at how much internal information about VSAM we in capacity and performance roles believe is common knowledge. Bereznay's article was full of excellent advice but, not to detract from it, he did overlook a major factor when defining CI freespace. That is, the number of bytes reserved as CI freespace should be large enough to hold at least one average length record:

Average record length = 630 bytes
CI size = 4096 @ 10% CI freespace
only 409 bytes are available for inserts.

Therefore:

CI freespace % = (# inserts per CI * average record length)/ CI size.

Additionally, specifying a fixed record length that is a multiple of 256 will waste almost one record in each CI due to the seven byte CI overhead for CIDF/RDF (CI Definition Field/ Record Definition Field) information. So subtract seven from your fixed record CI size, then divide that by the record length to determine the number of unusable bytes in a CI. Example:

(4096-7)/256 = 15 records/CI
with 249 unusable bytes left in each CI.

Contrast to a 255 byte fixed record which would leave only nine bytes unusable per CI!

David A. Stern II
BancBoston Mortgage Company
Jacksonville, FL

I enjoyed Frank Bereznay's article, "VSAM Specification and Tuning" (May/June 1988). The information provided was both interesting and useful. The discussion regarding the unused CA space as a result of poor key compression was of particular interest because this is an area where there is no direct warning and can be hard to spot.

There is one area in the article that requires further clarification. In trying to reproduce the poor key compression, Bereznay presents an AMS DEFINE CLUSTER with an index control interval size of 512 bytes. In reality VSAM will not allow a smaller index CI size than what would have been generated as a default. In this particular case the defaulted value would have been 2048 bytes. This is the size required by the index in order to hold 140 sequence set keys plus control information. There are ten 4K data records per track on a 3380 and as IMBED was specified, only 14 of the 15 tracks were used for data CIs. If NOIMBED had been specified, then 150 keys would have been required. In many cases VSAM is more

See Reader Forum *page 10*

JOBTRAC™

PRODUCTION CONTROL AUTOMATION

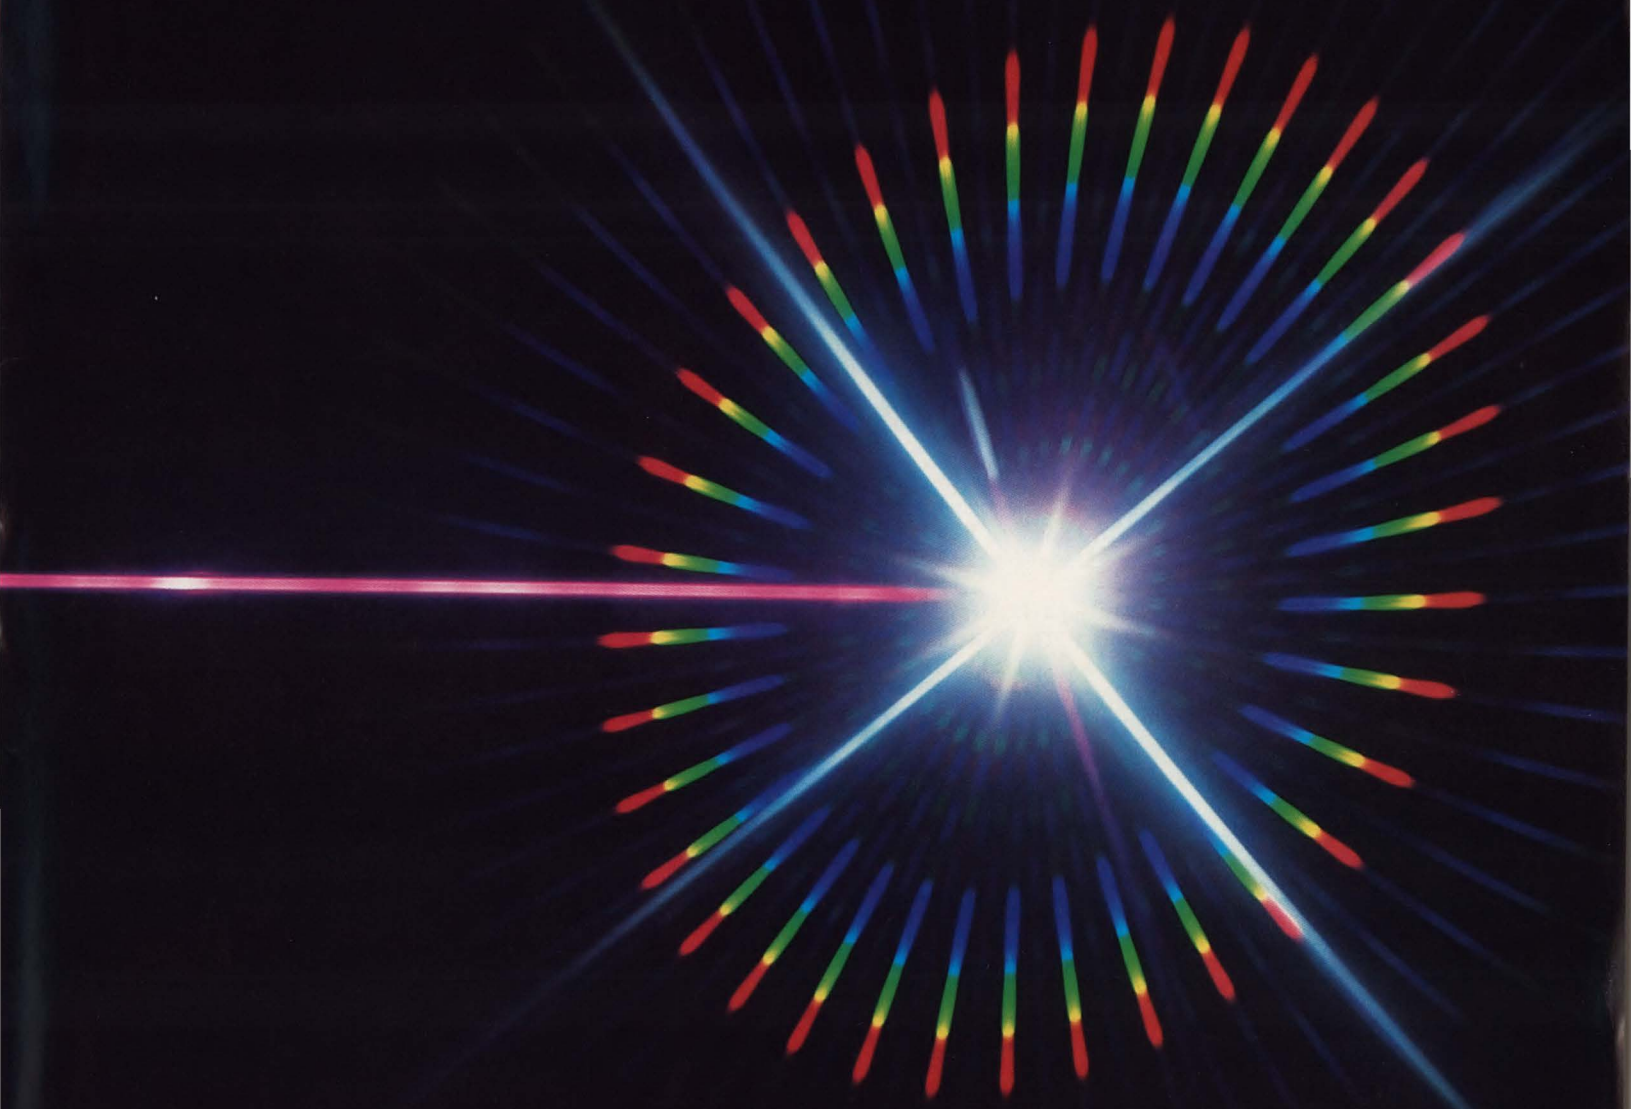

Today's Answer To Tonight's Problems.

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

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

CIRCLE #136 on Reader Service Card ▲

concerned with the number of keys to be accommodated in the sequence set record than the key length. The guidelines are:

# of Keys	Default Index CISZ
1-58	512
59-120	1024
121-248	2048
249-+	4096

In order to validate the assumption, I ran a similar cluster definition to the one contained in the magazine article. Even though 512 bytes was specified for the index CISZ, VSAM imposed an index CISZ of 2048 bytes. No warning message was sent to inform the user of the change. Only by evaluating a LISTCAT of the cluster can the user discover that a different index CISZ was generated. (Note: DOS/VSE is different in its assumptions as well as the new VSAM available with MVS/XA DFP V2.2.

Eugene S. Hudders
Multiple Computer Services, Inc.
San Juan, PR

MULTIPLE SESSION MANAGEMENT

I just finished reading the article in your July/August 1988 issue concerning VTAM multiple session management. In this article the author discussed software products that provide multiple sessions, while he wrote only three paragraphs covering hardware solutions.

While a hardware-based solution might not be right for everyone, it surely does not provide all the features that a software product provides. Because I work in a "True Blue" shop, I cannot address what products are available from other vendors, but I can address some of IBM hardware solutions.

For 3274 users that have model 41A, 41C and 41D controllers, there is a free RPQ available, 8X0002 Dual Logic RPQ. While the RPQ is free, you do need to add at least one control storage card to the 3274 to make use of the RPQ (feature code #3660 which has a purchase price of $1,550). With this RPQ, all CUT mode terminals attached to the control unit, except the terminal on port 0, can have two sessions.

For 3174 users that are running at microcode level A3.0 or greater, there is a feature built into the microcode called Multiple Logical Terminal support (MLT). With MLT support, CUT mode terminal users can have up to five sessions each. The exact number of sessions that can be supported on the 3174 is determined by the amount of storage configured on the 3174 and the type of terminal that is being attached. Using the 3178 terminal, it would be possible to provide up to 12 terminals, two sessions each, without having to buy additional storage. Depending on what other options are configured on the 3174, it could be possible to support more sessions all without additional charge.

For the present time, we have decided to use the hardware-based solution for multiple sessions. Some of the reasons for making this decision are:

- Performance — with a host-based solution, resources are used in the host for storing the screens for the inactive sessions using host storage. Also, all screens must flow thru the multiple sessions application using host cycles. From the network side, every time a terminal user switches sessions, the new screen that the user wants to overlay must be sent back up to the host so that any changes that were made are saved. This adds to network overhead. Using the 3X74 solution, the inactive screens are saved in the control storage of the control unit. Any session swapping that takes place is done by the 3X74. No host or network interaction takes place which in turn saves host cycles and network bandwidth.
- Reliability — if the host-based multiple sessions manager should fail, all users that were logged on thru the session manager would lose their sessions. With the 3X74 solution, if there were a problem with the control unit, the most that you would lose is 32 users.
- Usability — with the 3X74 solution the end user does not have to learn a new way to logon to the session manager and what commands to use to swap sessions or to establish new sessions. Using a 3174 to swap sessions, all the users have to know is to press the ALT and PA2 keys to cycle through their sessions.
- Problem Determination — by using the hardware solution, you eliminate the finger pointing that can take place when working on data stream problems. Is it the end user's application or is it the multiple session manager that is causing the data stream to go bad? Also, since most multiple session managers use a pool of VTAM applications to set up the end users application to the session manager connection, it becomes harder to try to figure out what LU name the end user is actually using to access something like CICS. All CICS knows is that it is talking to the session manager. With the 3X74 solution, the applications like CICS are still talking to a real honest-to-goodness terminal.

The bottom line is do not ignore some of the advantages that a hardware solution has over a software solution.

Richard Dougherty
Lumbermens Mutual Casualty Co.
Long Grove, IL

The article titled, "Multiple Session Management" (July/August 1988), was informative, well-organized and quite relevant to the market's interest in multiple session managers. Although the article covered all of the major points when evaluating multiple session managers, one crucial point was discussed in an incomplete manner, in my opinion.

I am referring to the discussion of the read buffer technique of screen management. Appropriately, the paragraph preceding the discussion of read buffer screen imaging identifies the three major side effects of session managers — additional memory requirements, additional CPU cycles and the user's mainframe path length. The read buffer (READBUF) technique is utilized by session managers specifically to mitigate against these adverse side effects. This discussion was a technical omission from the article and conceivably can mislead evaluators of session managers.

William C. Kolb
Westinghouse Systems Software
Pittsburgh, PA

I enjoyed the excellent article on VTAM session managers (July/August 1988), but was disappointed when I noticed that you inadvertently left MacKinney Systems off the VTAM session manager vendor list at the end of the article. We have the least expensive session manager that I know of and provide most of the features of the more expensive products.

VTAM/SWITCH allows users to switch from VTAM application to VTAM application (CICS, TSO, ICCF, IMS, TESTCICS, etc.) by pressing a PF/PA key. Multiple sessions of the same VTAM application are also allowed. VTAM/SWITCH has more than 100 users already, even though it was one of the last session managers to be introduced.

Rhonda Jenkins
MacKinney Systems
Springfield, MO

The July/August 1988 issue contained a list of VTAM session manager products. I noted that BIM's product, BIMWNDOW, which was one of the first in the market and remains one of the lowest priced, was omitted from the vendor list at the end of the article.

Bennett I. Moyle
B I Moyle Associates, Inc.
Minneapolis, MN

SNA's Logical Unit Type 6.2

LU 6.2 represents a departure from traditional SNA.

By David Peterson

The SNA Logical Unit (LU) type 6.2 has become the standard for connecting users in IBM's distributed processing environment. Communication takes place between transaction programs (TPs) that can be scheduled, executed and then terminated by the LU. Because of its design, LU 6.2 can be considered to be not only a program-to-program standard, but also a more general any-to-any standard that will continue to gain user acceptance.

Historical Perspective

Before LU type 6.2, SNA architecture defined the LU as a three-layer structure designed to support sessions with other LUs. These sessions are categorized by type according to how the two connected LUs use them to exchange data. Each session type uses a subset of the total range of SNA protocols that assures the most appropriate communication for a particular application. For example, LU session type 2 is designed for traffic between a terminal and host-based application consisting of the 3270 data stream. The session types are analogous to languages. For two LUs to communicate, they must use the same LU session type or speak the same language.

By defining multiple session types, IBM was able to address the different communication needs of its users. At the same time, however, this proliferation of session types created a diversity which in many cases had been an impediment to connectivity. To assure connectivity, a single standard within SNA for user-to-user communication was needed.

At about the same time, when the need for a communication standard was being addressed, the data processing world saw an explosive growth in processing power distributed outside the mainframe. CPU cycles became cheap and the use of personal computers and minicomputers increased dramatically. Any communication standard devised by IBM would

REPTRAC™

PRODUCTION CONTROL AUTOMATION

Report Distribution and Archival System

Real Time Report Capture • Parm Libraries • Eliminates JCL Changes
• Automatic Copy Distribution and Routing • Multi-Job Reporting • ISPF Interactive • No JES or MVS Modifications
• No Reruns Needed to Reproduce Any Report.

Bennett

Bennett Software, Inc.

1-800-JOBTRAC

P.O. Box 96694 • Houston, Texas 77213

CIRCLE #139 on Reader Service Card ▲

FIGURE 1

A comparison of the old LU-to-LU communication (a) with the new LU type 6.2 communication (b).

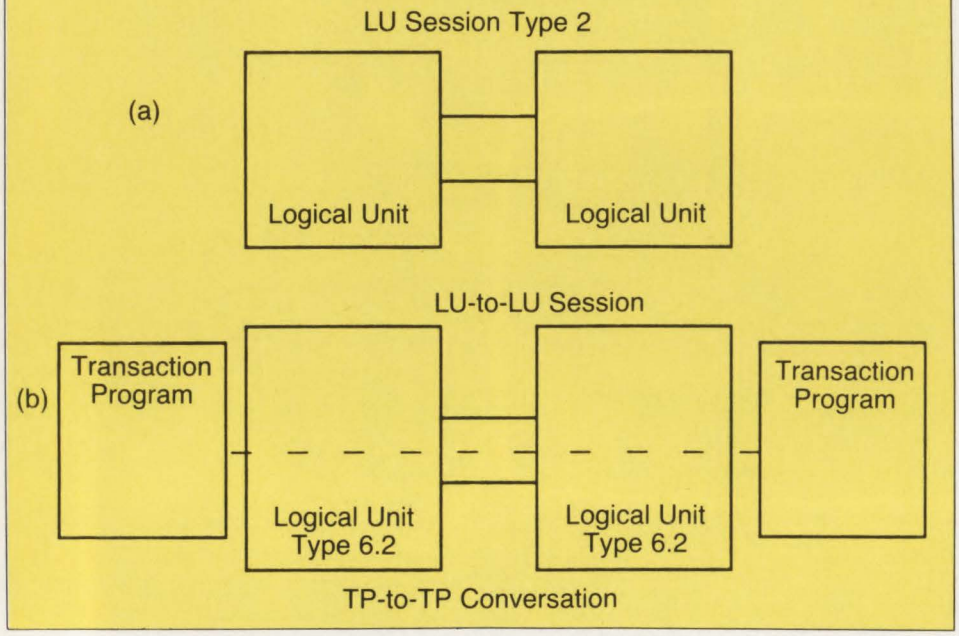

clearly have to address the requirements of the newly emerging distributed processing environment where data from different machines had to be exchanged. The standard that IBM introduced was Logical Unit type 6.2.

LU Structure Changes

LU 6.2 represents a departure from traditional SNA. Mechanisms for executing transaction programs have been woven into the LU's structure on top of the existing communication protocols. In this sense, LU 6.2 can be considered to be an enhancement to the old architecture.

User data is transferred over a temporary TP-to-TP connection called a *conversation*. A conversation is allocated by an LU at the request of one of its local TPs. This request causes a remote transaction program to be scheduled and executed. The TP-to-TP conversation traffic flows over the underlying SNA session that links the two LUs. With LU 6.2, the LU itself is categorized (that is, LU type 6.2), not the LU-to-LU session. See Figure 1 for an example. The SNA implementations, however, do not always make this distinction. That is, externally LU 6.2 is treated as another session type and not as a separate and uniquely classified LU.

Before a conversation between two TPs can take place, an SNA session must be established between the two LUs. During a session between two type 6.2 LUs, one LU acts as the Primary Logical Unit (PLU), the other as the Secondary Logical Unit (SLU). This structure exists with previous LU-to-LU sessions in which the PLU typically has more session control and recovery responsibility. However, the LU type 6.2 is designed to act as either the PLU or the SLU depending on its configuration in the network and the particular session. This design, including more symmetric recovery roles for the two LUs, diminishes to some extent the importance of the PLU and SLU hierarchy.

The PLU always sends the BIND RU to the SLU in order to create the session. The BIND is considered to be negotiable. That is, the SLU can either accept the session parameters in the BIND as they are or change them and return the BIND response to the PLU. The PLU can then accept or reject the new parameters. Certain optional functions that enhance the required base protocols can be selected by the two LUs through the BIND parameters.

A key feature of LU 6.2 that is based on the execution of TPs is the design of the Advanced Program-to-Program Communication (APPC) verbs. These verbs form a *protocol boundary* between the TP and the LU. A TP uses the APPC verbs to exchange data with a remote TP and to control the underlying conversation.

The verbs are specified generically by the architecture: their definitions are not dependent on any one language or system. Each language that supplies the verbs can use its own type of implementation that follows the generic model produced by IBM. In theory, there is a consistent model for these verbs, but in reality there is a diversity of implementations. IBM is now in the process of creating a single, consistent LU 6.2 programming interface based on the APPC verbs for its Systems Applications Architecture (SAA).

A system that implements LU 6.2 based on the generic APPC model can be classified in one of two ways. With an *open* implementation, the user can write his own TPs using the APPC verbs. A *closed* implementation does not allow such customization. Rather, it provides all the TPs necessary to perform the functions required by the specific system.

LU type 6.2 consists of four major SNA layers. The bottom two, Data Flow Control and Transmission Control, form the mechanism that sends and receives the user data over the SNA session. The third layer, Presentation Services (PS), is designed to manage the TPs local to the LU, including their interaction with the network as specified through the APPC verbs. The Transaction Services (TS) layer includes the architectural specifications for special transaction programs considered to be a part of the LU. These programs include LU service programs and other strategic architectures such as those used for the remote access of data. For example, two LUs can be engaged in multiple or parallel sessions. The Change Number of Sessions (CNOS) service transaction program assists in managing these sessions.

The end user of the LU, generally understood to be an application program, is positioned above Transaction Services, outside the LU. See Figure 2 for an illustration. The user-written TPs can use APPC verbs to send and receive data with other TPs over a conversation. Also, a special TP called the Control-Operator Transaction Program uses certain APPC verbs to control the operation of the LU.

Two Types of Conversations

The data flowing between two LUs conforms to the SNA General Data Stream (GDS) or to a user specified format. The GDS may be in one of several formats depending upon the two TPs exchanging the data. Each format, or subtype, of the

THERE IS A BETTER WAY TO MANAGE CICS: CICS MANAGER.

CICS MANAGER Version 2 is **the** better way, the new approach to CICS management. An approach that helps you avoid CICS crises and meet growing service level demands 24-hours a day, 7-days a week.

The push to 100% availability.
CICS MANAGER gives you the tools you need to achieve optimum availability and response time. Users have found that CICS MANAGER lets them maximize CICS performance, achieving greater than **99%** availability with sub-second response time. CICS MANAGER operates outside of CICS, so it is not affected by CICS performance problems. And your time is spent managing CICS – rather than reacting to crises.

More than a monitor. It's a manager.
CICS MANAGER lets you actively manage CICS performance:

- Warning screens alert you to impending problems.
- Informational data helps you find the cause.
- Extensive Action services let you execute effective solutions.

You can tune multiple, **even remote**, CICS regions from a single terminal. Switching across regions, machines and data centers is as simple as changing the name of the CICS region on your screen. Even novice users can identify and solve problems within a single session.

CICS MANAGER puts you in control.
With detailed on-line monitoring. Instant replay of recent data for system tuning. A unique Performance Reporting Language that lets you design historical reports for efficient planning. And CICS MANAGER works in concert with other Boole & Babbage performance tools, so you can *get the most* from your entire system.

See how it can work for you.
For a free demonstration of CICS MANAGER, call Marty Johnson today. In California: 800-624-5566. Outside California: 800-822-6653.

Boole & Babbage

International sales and support provided through The European Software Company and a worldwide distribution network.

CIRCLE #160 on Reader Service Card ▲

data stream consists of its own structured fields and has a unique identification (GDSID).

Data is transferred between TPs as logical records. Each record begins with a two-byte length field (LL). An APPC SEND_DATA verb causes one or multiple logical records to be passed to the local LU for transmission. If the data is formatted according to the GDS, the first record will include a two-byte GDSID identifier after the length field that specifies the format of the data.

Two types of TP-to-TP connections are defined according to how the TP accesses the data. During a *basic* conversation, the TP must handle the actual logical data record, including the two-byte LL field when sending and receiving data. It can use a GDSID to identify the format of the following data. This program issues the basic form of the APPC verbs. A TP involved in a *mapped* conversation deals only with the user data, possibly in an unstructured form. In this case, the Presentation Services layer of the LU inserts the LL field and can translate or map the user data into a structured format as specified by the TP. The GDSID is also inserted by PS. The verbs used in this case are called the mapped APPC verbs.

TPs engaged in basic conversations are usually supplied by IBM or follow an IBM design, while those that use mapped conversations are typically user-written. The IBM-designed TPs are considered to be within the Transaction Services layer and therefore part of the LU (refer to Figure 2). Because of the additional PS processing involved during a mapped conversation, IBM defines a separate basic and mapped TP protocol boundary. However, both the IBM-designated programs and user TPs interact with the Presentation Services layer of the LU. Therefore, only one protocol boundary is shown in the figure between the TP and Presentation Services.

Every TP is identified by a name (TPN). IBM has reserved certain names that begin with a character of less than EBCDIC hex '40' for SNA-defined TPs. The local or source TP issues the APPC verb ALLOCATE in order to schedule the execution of a remote or target program. The ALLOCATE verb causes the local LU to send a request containing the name of the target TP to the remote LU. The remote LU then creates the mechanisms necessary to execute and terminate the selected TP.

FIGURE 2

The layers of Logical Unit Type 6.2 reveal the TP protocol boundary.

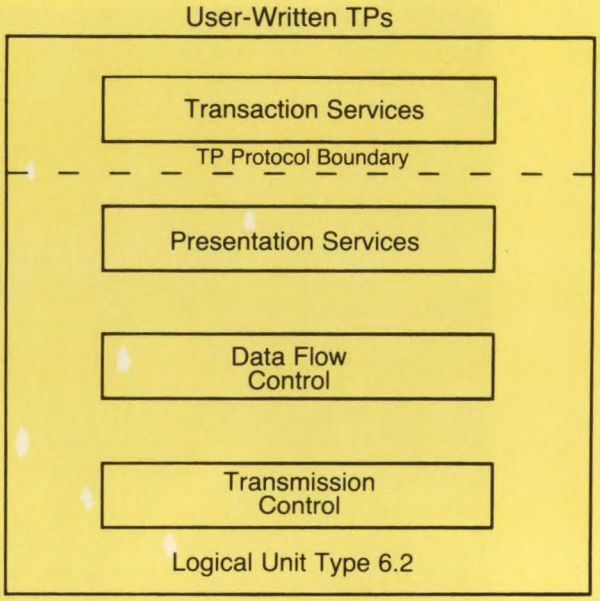

Half-Session Protocols

The SNA layers Data Flow Control (DFC) and Transmission Control (TC) form the basis of the mechanism that sends and receives data over the SNA session. This half-session protocol mechanism or machine allows TP-to-TP data to be exchanged between LUs.

LU type 6.2 uses the DFC send-receive mode protocol called half-duplex flip-flop with brackets. This protocol is also used by LU session types other than LU type 6.2 The bracket defines a conversation; it consists of a sequence of bi-directional flows between the LUs.

After the BIND, the session is within a bracket and there is no contention. At this time, the PLU is always in send state and the SLU is in receive state. The PLU can use the session for a TP-to-TP conversation or immediately end the bracket causing the session to enter bracket contention state.

Once the session is in contention state, either LU may begin a bracket. However, a problem arises when both LUs attempt to initiate a bracket at the same time. To resolve this situation, one of the LUs is designated the first speaker (FSP) and the other the bidder. The FSP can begin a bracket at any time, so it always wins in a contention situation. The bidder must ask for and receive permission from the FSP to begin a bracket. The FSP and bidder roles are negotiated during the BIND process.

The APPC verb ALLOCATE is used by a TP to start a conversation with another TP. The type 5 Function Management Header (FMH-5) and the begin bracket (BB) flag flow from the local to the remote LU in order to request the conversation.

If the TP is connected (via the APPC verbs) to the session's FSP, the FMH-5 and BB are sent immediately to begin the conversation. When a TP is connected to the LU acting as bidder, the LU must bid to begin the bracket. This bid is an explicit request sent before the FMH-5 consisting of the LUSTAT RU and BB flag. Alternately, the bid may be made implicitly. In this case, the LU immediately sends the FMH-5 and BB with the possibility that the request might be rejected.

Once a bracket has been initiated, the two TPs can exchange data using the APPC verbs that cause the LUs to alternate between send and receive states. The Change-of-Direction Indicator is used to signal a flip-flop between send and receive states.

The APPC verb DEALLOCATE is used by one of the TPs to request that the conversation be terminated. Termination sets the conditional end bracket (CEB) flag on the last LU-to-LU request unit chain element.

If...

One product and one hour could double VSAM performance

FIGURE 3

A simple TP-to-TP conversation using LU type 6.2.

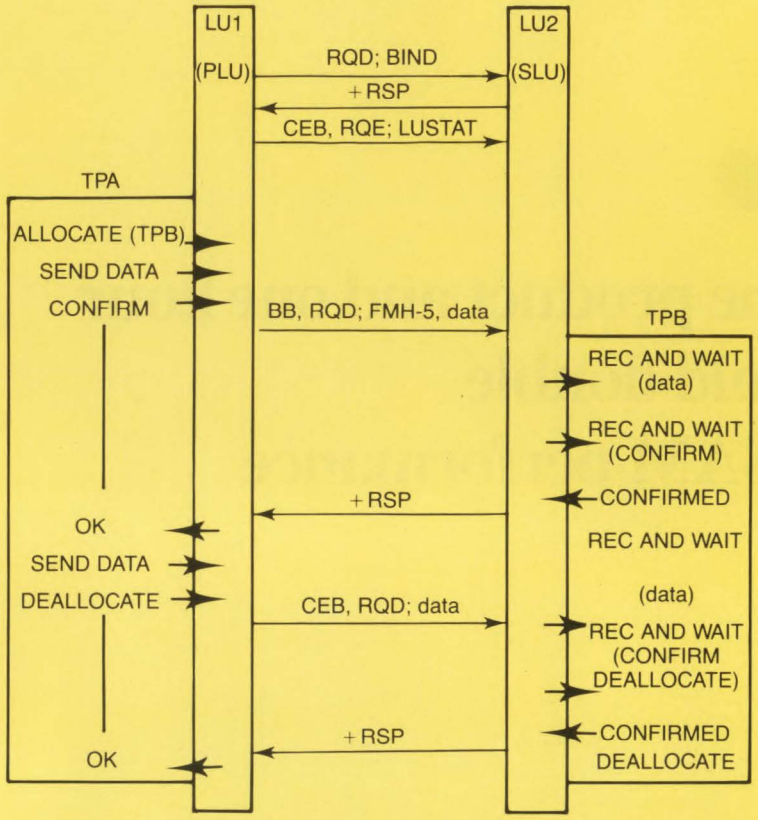

All LU-to-LU flows use single chain elements. BB: begin bracket; CEB: conditional end bracket; FMH-5: function management header, type 5; RQD: request definite response; RQE: request exception response; RSP: response (either + or −).

Both definite and exception response protocols can be used by LU 6.2. One type of response is reserved by the base LU for use over the underlying SNA session in certain cases (DR1), while others are used by the two TPs for their conversation.

Another DFC protocol used by LU 6.2 is chaining. Chaining allows an LU, once it receives data from a local TP, to break up large messages into pieces small enough for the network and the two LUs to handle. The messages are then sent to the remote LU where they are reassembled and presented to the target TP.

LU 6.2 can also use the TC-layer protocols of pacing and data encryption. Pacing controls the rate of traffic that flows over a session. Data encryption can be used for sensitive data.

Sample Conversation

Figure 3 shows a simplified example of an LU 6.2 conversation. The BIND is sent to the SLU in response to a request or some external event, perhaps system start-up. After the positive response to the BIND is received by LU1, the session falls by default within a bracket. The PLU is in the send state and the SLU is in the receive state. To terminate the bracket and enter contention state, the PLU sends the conditional end bracket flag with the LUSTAT request unit (LUSTAT has no significance and acts as a null RU). After some time, TPA is scheduled at LU1, perhaps as the result of an operator command. TPA begins execution by requesting a session with TPB through the ALLOCATE verb. The FMH-5 is buffered along with data passed by the SEND command until the CONFIRM verb is issued by TPA. The data is then sent to LU2 along with the request for a definite response. Finally, TPA requests that the conversation be ended with the DEALLOCATE verb. After receiving the definite response from LU2, TPA is terminated.

Future of LU type 6.2

In the past, IBM's seal of approval has helped to establish particular systems and architectures. With LU 6.2, the process has been repeated. IBM has designated LU type 6.2 as its strategic communication protocol. Since it is also included under the Systems Application Architecture umbrella, IBM assures it of long-term user acceptance. Although current LU 6.2 offerings are limited, IBM will continue to expand its implementations. For example, VTAM version 3.2 will have an Assembler language access to LU 6.2 services; it will be even easier to build new application systems that use LU 6.2. Also, the completion and acceptance of OS/2 Extended Edition will create a large pool of PS/2 microcomputers linked in some way to a host-based transaction processing system over LU 6.2.

We can expect to see existing products enhanced to use the new LU protocols. For example, the 3174 control unit could be upgraded by IBM to allow existing 3270 type terminals to use LU 6.2 protocols instead of the current LU session type 2.

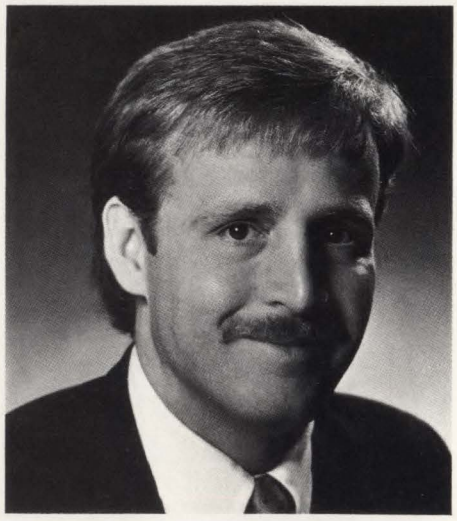

ABOUT THE AUTHOR

David Peterson is a member of the R&D staff at Candle Corporation where he is currently involved in the area of IBM communication technology. Before joining Candle, Peterson worked for an insurance company and a major manufacturer providing technical systems and application programming support. In addition to VTAM and Netview, he has worked extensively with MVS, VM, CICS and computer graphics. Candle Corporation, 1999 Bundy Drive, Los Angeles, CA 90025; (213) 207-1400.

Then...

Introducing HYPER-BUF, the VSAM buffer allocation expert

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

Other VSAM products from Goal Systems, the Original VSAM Experts

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

To Find Quality Call 800-848-4640.

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

Goal Systems

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

CIRCLE #79 on Reader Service Card ▲

Does your staff spend more time waiting . . . when they could be working?

Then you need DOCS!

By operating your VSE system from just one console, you're not only wasting time . . . you're wasting money! Because instead of working, your staff is whiling away the hours waiting to check on their jobs.

But with DOCS, Display Operator Console Support from SMARTECH, you can operate your VSE or MVT system *from up to 48 true consoles — local or remote.* So your staff can monitor jobs from their own desks, and you get a more productive computer room. That means more work is completed in less time . . . and time is money!

DOCS also saves you time and money in many other ways. It tells you instantly which jobs are running, which aren't *and* why! And DOCS shares your terminals with your other on-line software — *with instant hot-key switching.*

DOCS™

Best of all, DOCS takes just 30 minutes to install and requires no assemblies or customization. And you'll be supported by the industry's best technical support who are only a free phone call away.

Try DOCS free for 30 days!

See for yourself how DOCS can save you money! Try it free for 30 days. Or for more information . . . and a free reprint of this ad call . . .

David Shepherd at 1-800-53-SMART
Outside the U.S., call (214) 630-8324

SMARTECH SYSTEMS, INC.

Turning high technology into *SMART TECH*nology.℠

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

CIRCLE #65 on Reader Service Card ▲

Pete Clark's

VSE Forum

Question: If Pete Clark is truly interested in what to do for an encore, how about more VSE partitions? Four more should be easy for Pete! Sure would be great to get rid of VM's overhead requirement when all we really need is more partitions.

Stephen P. Fry
Allen County

Clark: I was contacted by a user early in July who was very interested in altering the VSE code to support additional partitions. We had a very productive discussion on the possible problems and the techniques with adding partitions. The user (who shall at this time be called John) assured me that he believed that he could resolve the issue and would be willing to share the code with anyone wishing to use it.

To support additional partitions within the confines of today's VSE will dictate that ICCF not be available on the system. We understand that this may be a problem for some users who have a requirement for ICCF.

John is attempting to have the code working, tested and available this fall. I certainly wish him good luck, Godspeed and happy coding. We are looking forward to testing his patch and as more information becomes available, MAINFRAME JOURNAL will keep you informed.

Question: Please consider an article on a native VSE/SP-to-VM (multiple VSE guests) conversion. Since it appears VSE is here to stay, this article would be quite useful.

Rich Szabo
GCC Inc.

Clark: I really do not like running a production VSE/SP under VM because of the performance implications, but I certainly recognize the value of running applications testing, systems testing and a development VSE/SP guest under VM.

VM is an excellent way to maximize hardware resources and yes, your first question would make an excellent article. I'm sorry that I do not have the time and space to do your question justice now. Perhaps later MAINFRAME JOURNAL will do an article. Yes, I do believe that VSE is definitely here to stay and as evidence of this have you noticed:

- The rapid acceptance, announcement, delivery date for the address space patch
- The VSE/SP 3.2 IBM announcement
- VSE statement of direction contained in the VSE/SP 3.2 announcement
- The IBM SAA VSE press release
- The IBM VSE office product announcements
- The PR/SM support of VSE as a full participating guest.

It is my perception that VSE/SP has a bright future as the IBM midrange 370 operating system.

Question: Why not expand on the VSE 90M virtual storage extension? For example, how would one use the address spaces if most of your area is sharable? How do you use a second CPU for the spooling (POWER) system?

John Murray
Keyes Fibre

Clark: Assuming that your shared area is 8M of the 16M and perhaps you require a 12, 16 or 24M for CICS, then when utilizing CICS/VSE release 1.7, consider running several CICS partitions utilizing MRO facilities for communication between the CICS partitions. There are several ways to divide CICS up. Consider a terminal-owning partition, a file-owning partition, several separate application-owning partitions and/or a partition containing any combination of the preceding areas.

We currently utilize the POWER shared spool facility with two VSE machines participating in sharing the spool files. CPU 1 is the production VSE and contains no typical input or output facilities for RJE connections. The object of this is threefold: to reduce the virtual storage required in the production POWER partition (CPU 1, note that it is a shared partition and reducing shared returns storage in all address spaces); to remove slow speed I/O devices from the production CPU; and to use some of the virtual saved by item 1 to enlarge the POWER data block size to improve POWER performance and reduce I/O.

Yes, we understand that we have introduced cross system DASD sharing by using this type of arrangement, so to help reduce the impact of cross system sharing we have installed EXTEND/VSE, a lock file management facility from Goal Systems (Columbus, OH).

Question: I would like some information about the "double paging" problem when attempting to use VAE (especially with "the patch" from Pete Clark) under VM. What can we VSE-under-VM users do?

Mike Petonic
Cycare Systems, Inc.

Clark: First the "double paging" considerations are not changed by implementation of the patch. All rules, regulations and caveats definitely apply. Second, if you want the best VSE performance then you must run VSE native. When running VSE/SP in virtual address extended mode (VAE) under VM, the best performance will be obtained by utilizing the virtual = real (V = R) facility of VM. Unfortunately, this option is limited to only one VM guest and requires dedicating real memory to the guest, limiting real memory resources to other users.

Running VSE/SP in VAE mode under VM with virtual = virtual (V = V) will invoke "double paging" and "double CCW translation" that translates into additional CPU overhead and addtional I/O. If you have excess CPU and I/O capacity, the impact of running V = V may be acceptable. However, if you are already running high CPU and I/O utilization, VAE V = V will not provide acceptable performance.

We currently run one VSE/SP VAE native on a 3083-B and two VSE/SP VAE V = V under VM on a 3083-E. Both CPUs have 16 MB real memory. Performance of the native VSE is exceptional and the performance of the two V = V guests in our environment (with our load) is also extremely good. If you have the resources, I would certainly try VAE V = V to see if the performance is acceptable.

Regarding the patch, the only real impact is the increase in virtual storage available to the guest. This may result in additional paging as you increase your use of additional address spaces and additional virtual storage. Of course, at some point you will over balance the virtual/real memory ratio and paging will become a problem. While it appears that this is certainly installation dependent, with 16M of real memory this over balance seems to be typically at the five or six address space point or between 50 to 70M of virtual. It is my contention that paging is a resource option not normally desirable when performance is the primary requirement; paging should always be minimized. We currently average a page per second on our production VSE and at six pages per second we experience performance problems. Paging consumes resources in two ways: increased I/O and increased CPU consumption.

ABOUT THE AUTHOR

Pete Clark has been in data processing for 25 years, the last 11 with Olan Mills, Chattanooga, TN. Before that he was in technical support and was a data processing teacher at Chattanooga State Technical Institute. Clark has been working many years to extend the limits of VSE. His latest effort to enhance VSE extended its capability well beyond Release 3.1.

Computer Associates introduces total data center automation for VSE.

Production Control

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

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

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

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

Call Dana Williams today:
800-645-3003 (Ext. 0000).

711 Stewart Ave., Garden City, N.Y. 11530-4787

Storage and Resource Management

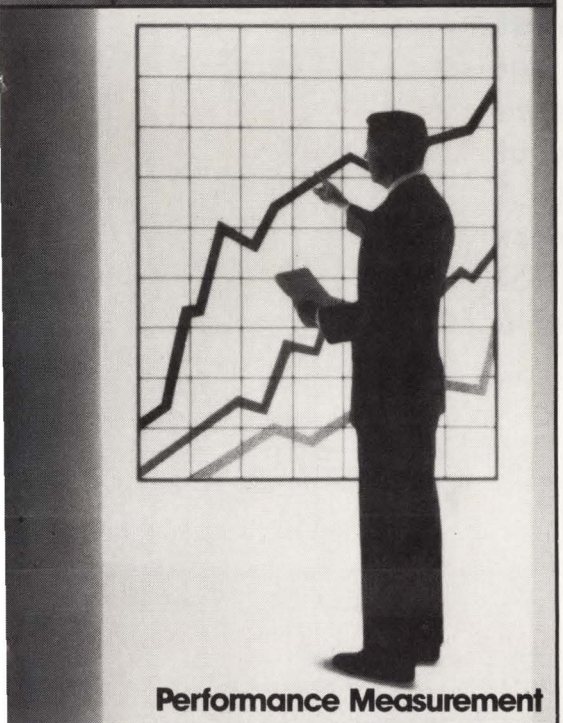

Performance Measurement

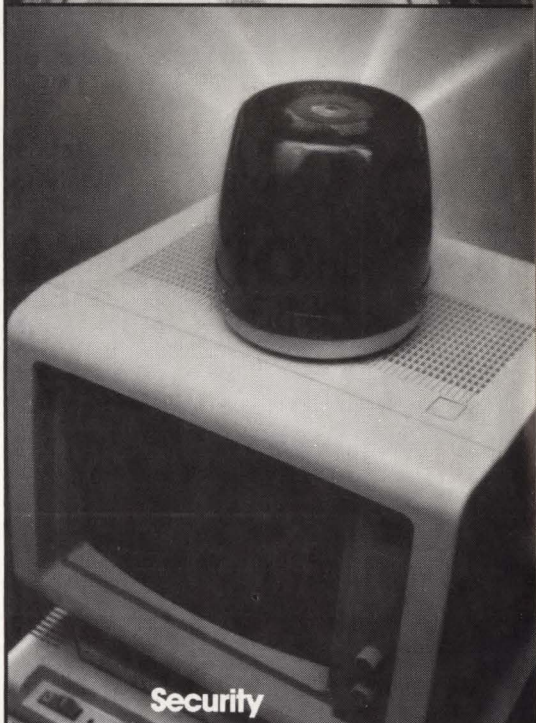

Security

COMPUTER® ASSOCIATES
Software superior by design.

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

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

CIRCLE #105 on Reader Service Card ▲

Key Compression... Is It Ruining Your DASD Utilization?

By David Martin

Tuning VSAM clusters for optimum DASD use is a complex and time consuming task. Many factors should be considered when tuning including DASD type and the type of processing. Dataset structure, data and index control interval (CI) size, freespace and record size are other important factors to consider when faced with the problem of limited DASD space. All of these factors play an important part in tuning. One small factor, however, that can destroy all of your tuning efforts is *key compression*. Understanding what key compression is and how it works can save DASD space that is mysteriously disappearing.

VSAM and Key Compression

One of the first questions to answer when defining the record layout for a VSAM cluster is, "What will the key consist of?" The key is a field in the record that serves as its unique identifier. The characters or digits that make up the key determine how VSAM compresses it.

FIGURE 1

Key Compression Table (VSAM 3)

KEY LENGTH	IDCAMS AVERAGE COMPRESSION VALUE
KEY LENGTH < 10	LENGTH OF KEY
10 < = 30	10
30 < = 64	KEY LENGTH / 3
64 < KEY LENGTH	25

The key-sequence dataset (KSDS) is one of three types of VSAM clusters and consists of a data component and an index component. The data component consists of data CIs while the index component consists of index CIs. When a KSDS is loaded, records are sequentially written until each data CI becomes full. The highest key in each data CI is then compressed, if possible, and put into the index CI. If the key is constructed in such a way that VSAM cannot compress it, the index CI can fill up before the data control area (CA) resulting in unused data CIs.

Among other variables, VSAM uses an *average key compression factor* to determine a valid index CI size. The averages can change from PTF level to PTF level. Figure 1 shows general key compression averages.

F I G U R E 2

Relationship of Index CI with Key Compression

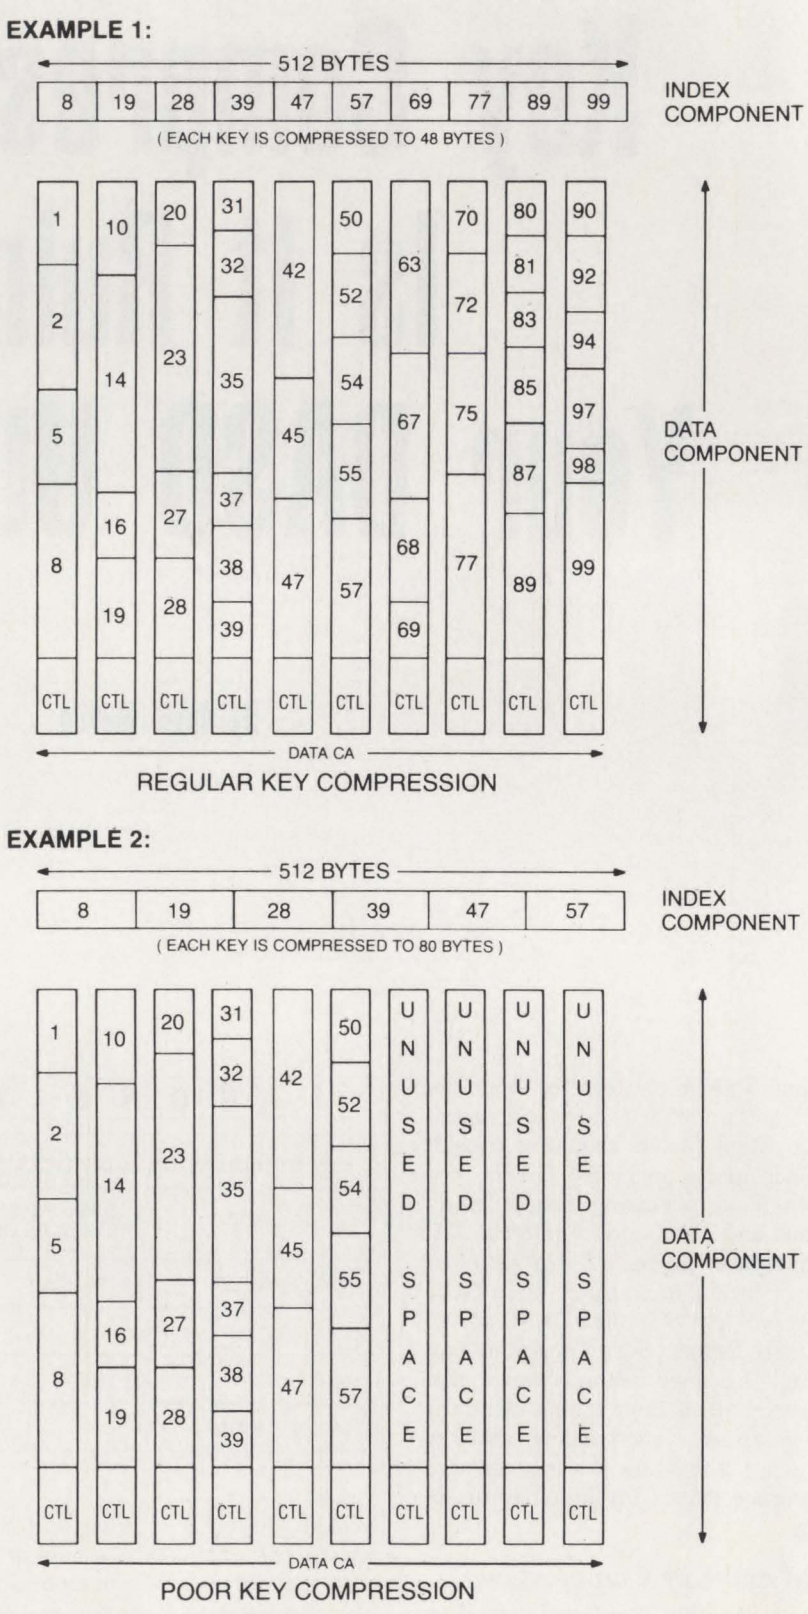

AT $1,595 THE MAJOR DIFFERENCE BETWEEN OUR 3270 PRINTER AND THEIRS.

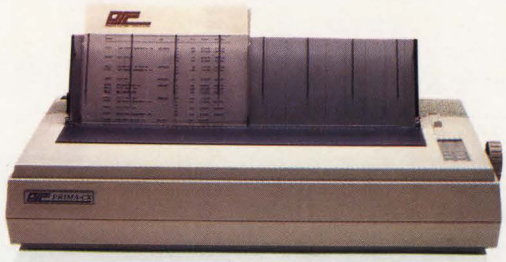

OURS EXISTS.

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.

PSC 1 YEAR SERVICE

PSC

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 ▲

FIGURE 3

Compressed Key, F and L Fields in Index CI

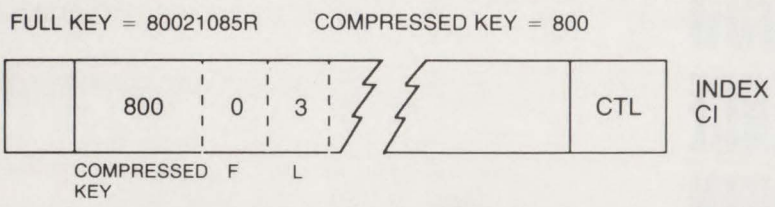

F = FRONT KEY COMPRESSION LENGTH
L = COMPRESS KEY LENGTH

(SEE FIRST DATA CI OF FIGURE 4 FOR KEY REFERENCE.)

The problem with the averages shown occurs when the keys in a cluster do not compress to what the average is. Since VSAM uses the average key compression length instead of the actual key compression length in the algorithms to determine a valid index CI size, VSAM can determine that the index CI size is large enough, even though there may not actually be enough space to contain all of the keys needed to index the associated data CA. Figure 2 contains two examples that illustrate the result of key compression. The first example shows the result of normal key compression and the second example shows the result of poor key compression. The key length is 90 bytes in the examples. Each data CA contains 10 data CIs so the index CI must have enough space to contain 10 compressed keys.

Many factors should be considered when tuning . . . One small factor, however, that can destroy all of your tuning efforts is key compression.

Understanding key compression can save DASD space that is mysteriously disappearing.

David Martin

In the first example in Figure 2, VSAM uses 25 as the average key compression length. Because of the contents, the key compresses to 48 bytes. Even though the key does not compress to the average, there is adequate room in the index CI for all ten compressed keys.

In the second example in Figure 2, the key compresses to only 80 bytes. VSAM still uses 25 as the average key compression length and determines that the index CI size is large enough to contain all 10 compressed keys. Because of poor key compression, only six of the 10 compressed keys from the data CA fit into the index CI. This results in 40 percent of

CICS
SPOOL DISPLAY and PRINT

CPMS®

■ SPOOL PRINTING
- ► Manual or automatic selection
- ► Specific reports selection
- ► Full FCB forms control
- ► Complete recovery and restarting
- ► Secured CPMS printer control

■ SDSF-LIKE SPOOL DISPLAY
- ► Entire job or specific report
- ► Search for character string
- ► 4-way scrolling
- ► Special capability to view: condition code summary, LOG, JCL, MSG, end of report, system log

■ JOB MANAGEMENT
- ► Display active jobs
- ► Route to system, RJE or CICS printers
- ► Hold, release, cancel or purge jobs
- ► JES queue status displays
- ► Full support for large screen devices
- ► Secured JES2 printer display/control

- Full security and security exits
- Menu driven with assignable PF keys
- MVS, MVS/XA with CICS
- Over 450 customers worldwide
- $9,500

H&W®

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

FIGURE 4

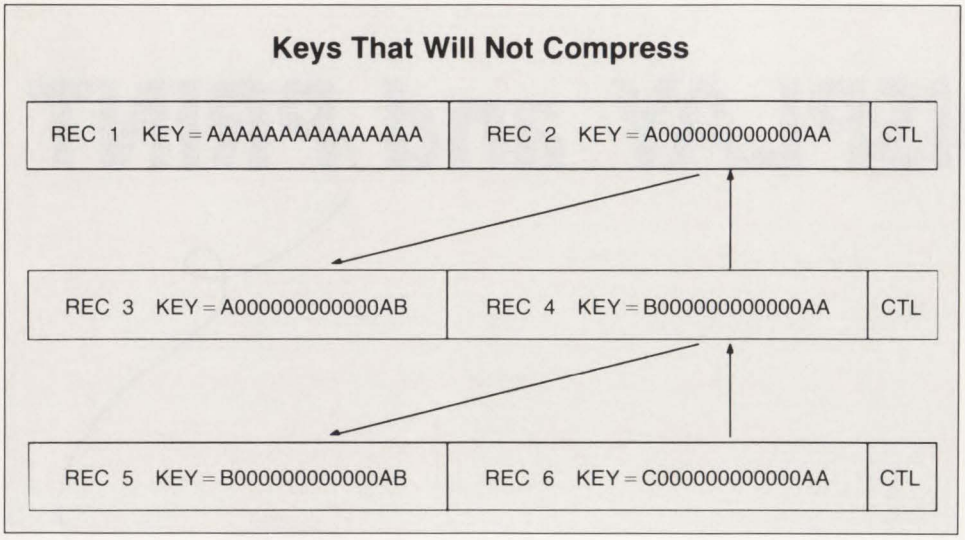

FIGURE 5

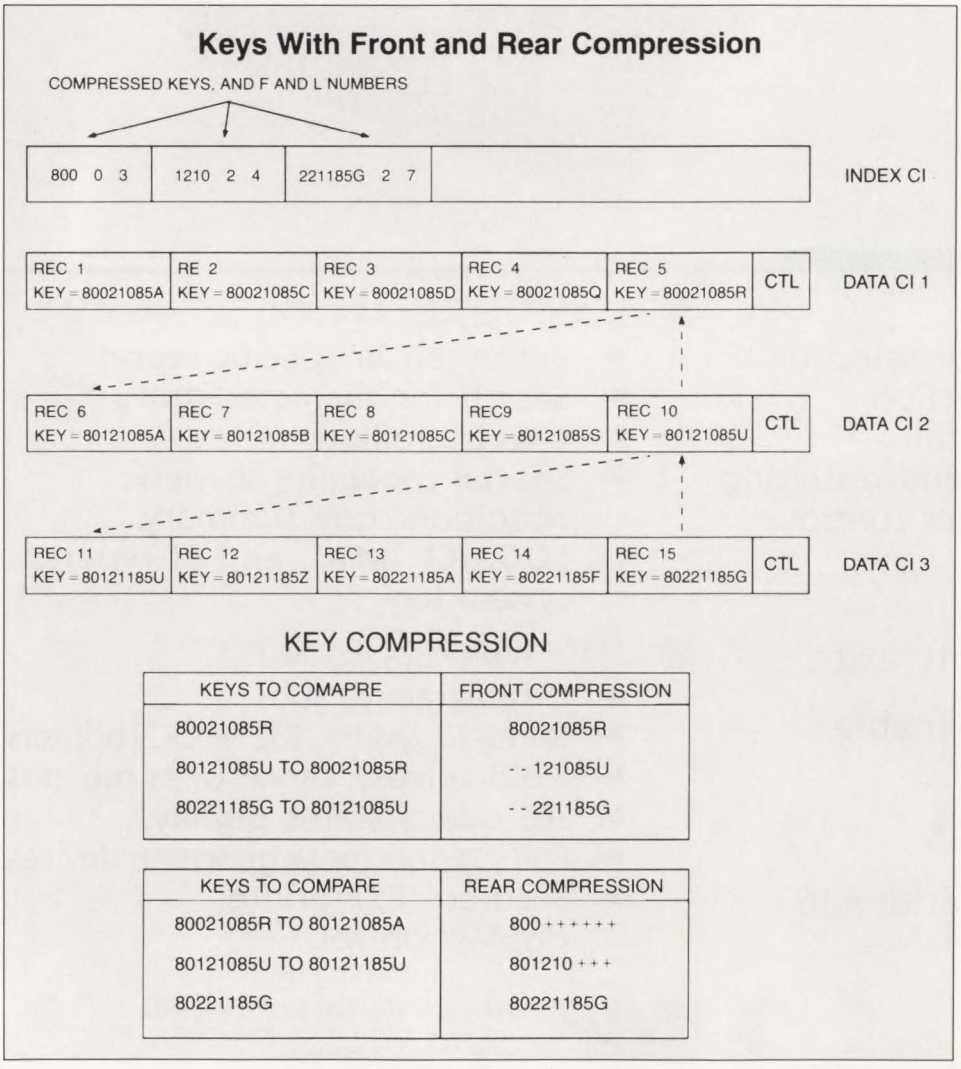

KEY COMPRESSION

KEYS TO COMAPRE	FRONT COMPRESSION
80021085R	80021085R
80121085U TO 80021085R	- - 121085U
80221185G TO 80121085U	- - 221185G

KEYS TO COMPARE	REAR COMPRESSION
80021085R TO 80121085A	800 + + + + + +
80121085U TO 80121185U	801210 + + +
80221185G	80221185G

each CA being left unused. Increasing the index CI size to 1,024 bytes would correct this problem of unusable space.

How Key Compression Works

Understanding how key compression works can help in constructing record keys that compress optimally.

VSAM begins by locating the highest key in each control interval and comparing it with the highest key in the previous control interval. The comparison starts with the leftmost byte of the key, compares in a forward direction and ends with the first unequal value. Identical values from the left are removed. This process is called *front key compression*.

VSAM then takes the highest key in each control interval and compares it to the lowest key in the next control interval. All values after the first unequal value are removed. This is called *rear key compression*.

Once the uncompressed key has gone through both key compression processes, the compressed key is put into the associated index CI with two other fields, the *F* and *L* fields. The F field contains the number of front key compressed bytes. The L field contains the length of the compressed key.

Figure 3 shows a diagram of an index CI with a compressed key. It should be noted that the highest key in the first data CI of the first data CA does not front compress and the highest key in the last data CI of the last data CA does not rear compress.

Figure 4 illustrates keys that cannot be compressed. The leftmost byte of the high key in CI 2 is greater than the leftmost byte of the high key in CI 1, resulting in no front compression.

The first unequal byte in the comparison between the keys in record two and record three is the last byte, resulting in no rear compression.

VSAM still uses ten as the average key compression length because the key length of the cluster is 15 bytes, even though the key cannot be compressed.

In Figure 5, there are five records per each data CI. The key length is nine bytes. The keys are constructed so that front and rear compression takes place. Notice the first key does not front compress and the last key does not rear compress.

Once again in Figure 6, the keys are constructed so that front key compression takes place. Notice that the third data CI contains a key that fully compresses.

IAM vs VSAM

The Incredible Shrinking Machine

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

IAM FILE STRUCTURE

IAM SAVES 20 to 40% DASD SPACE

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

DATA COMPRESSION

SAVES AN ADDITIONAL 20 to 50% DASD SPACE

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

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

RELEASE

AUTOMATIC RELEASE OF UNUSED SPACE

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

TRANSPARENT

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

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

SMF ANALYSIS

VSAM SIZE REPORT

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

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

Call for a Free No Obligation 90 Day Trial

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

INNOVATION DATA PROCESSING

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

CIRCLE #97 on Reader Service Card ▲

FIGURE 6

Front and Rear Key Compression

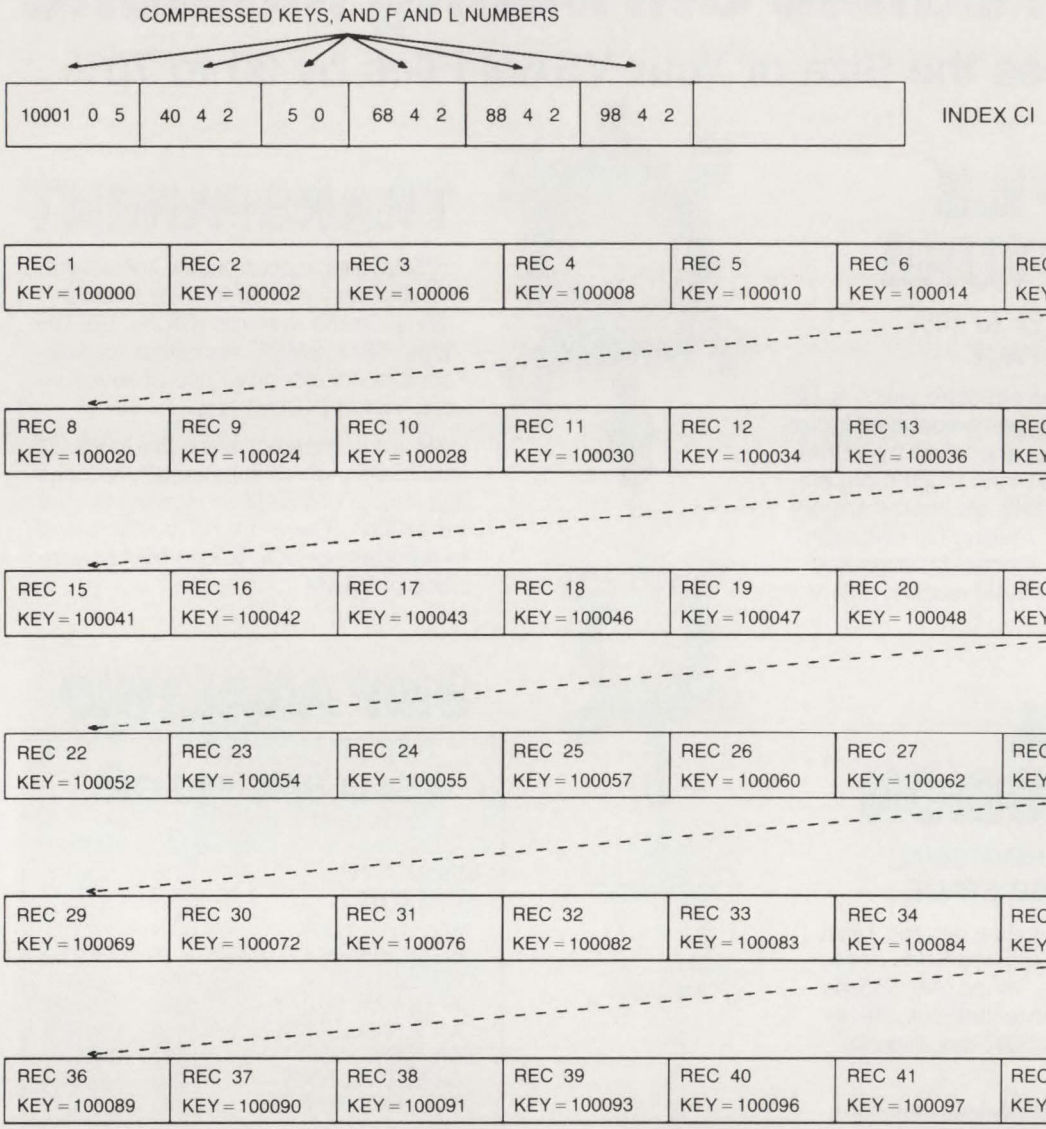

KEY COMPRESSION

KEYS TO COMPARE	FRONT COMPRESSION
100018	100018
100040 TO 100018	- - - - 40
100049 TO 100040	- - - - - 9
100068 TO 100049	- - - - 68
100088 TO 100068	- - - - 88
100098 TO 100088	- - - - 98

KEYS TO COMPARE	REAR COMPRESSION
100018 TO 100020	10001 +
100040 TO 100041	100040
100049 TO 100052	10004 +
100068 TO 100069	100068
100088 TO 100089	100088
100098	100098

FIGURE 7

Reconstructed Keys

FULL KEY	COMPRESSED KEY	FRONT COMPRESSION LENGTH (F)	COMPRESSED KEY LENGTH (L)	RECONSTRUCTED KEY
100018	10001+	0	5	10001F
100040	----40	4	2	100040
100049	-----+	5	0	10004F
100068	----68	4	2	100068
100088	----88	4	2	100088
100098	----98	4	2	100098

ABOUT THE AUTHOR

David Martin is an experienced product developer in the Data Center Management Division at Goal Systems and the original author of the MASTERCAT product, an on-line VSAM display facility. Martin has been directly involved in development and support of VSAM products for Goal Systems for more than five years. Goal Systems International, 7965 N. High St. Columbus, OH 43235, (614) 888-1775.

Key Expansion

To expand a key, VSAM takes the uncompressed values from the previous uncompressed key and rebuilds the front of the compressed key. VSAM then substitutes the rear compressed values with high values (X'FF'). The number of front key compressed bytes and the compressed key length are used in reconstructing the keys. Figure 7 shows what the compressed keys look like after they have been reconstructed.

Key Compression Considerations

It is sometimes difficult to identify poor key compression conditions. The compression can change with record additions, updates and deletions. It can also change with cluster reorganization. One symptom that is very noticable, however, is a cluster that requires much more space than calculated or that runs out of extents before anticipated. Another symptom is many CA splits occurring on record additions or updates. Monitor both symptoms with a catalog listing.

It is also difficult to establish a rule of thumb for building keys. Here, however, are a few considerations. Multiple field keys, called *complex keys,* can compress poorly, especially when changes occur in the high and low ends of the key. Front compression is best when the highest key of each data CI has the same leading characters as the previous high key. Rear compression is best when adjacent keys have large variations in the right most characters. Figure 6 illustrates good front and rear compression.

By disregarding the concept of key compression, the DASD space that has been recovered because of other types of tuning can be lost again if the keys do not compress well. Take the time to see if key compression is a problem with some of your datasets. Simply changing an index CI size can save critical DASD space. You may discover you did not need that extra 3380 after all . . .

3090 PR/SM Feature Provides Configuration Flexibility

By H. Pat Artis and Alan Sherkow

The IBM 3090 Processor Resource/Systems Manager (PR/SM) feature introduces a new multi-image capability for 3090-E models. The 3090-E models now have three operating modes: S/370, ESA/370 (which supports 370-XA and ESA/370 operating systems) and Logically Partitioned (LPAR) mode.

This provides flexible partitioning into as many as four logical partitions. On 3090-E models that can operate physically partitioned (3090 Model 280E, 400E, 500E or 600E), each physical partition can have four logical partitions for a maximum of eight partitions on these models. When the LPAR mode is chosen, the operator can define the resources that are to be allocated to each logical partition. The resources that are distributed between the LPARs are central processors (CPs), channel paths, central storage and expanded storage.

The PR/SM feature is a set of hardware and microcode features that allow non-physical partitioning or logical partitioning. Physically partitionable processors (3084-Qs and some 3090 models) have the ability to partition along physical power boundaries. These systems, when partitioned, provide two identical halves that are almost fully independent of each other — only separate processors have greater isolation. When logically partitioned, there is less isolation between the partitions than the physically-partitioned modes. Logical partitioning is available on the entire 3090-E family providing more flexibility than physical partitioning at the cost of isolation.

Most resources can be reconfigured without a 3090 power-on-reset. Reconfiguration of CPs or memory will require shutting down one or more LPARs. Once a logical partition has been activated, a supported operating system can be IPLed into that logical partition. These include:

- MVS: MVS/SP 3 (that supports ESA/370), MVS/SP 2 and MVS/SP 1 Release 3.5 and Release 3.6
- VM: VM/XA System Product Release 1 and Release 2, VS/SP HPO Release 5 and VS/SP Release 5
- VSE: VSE/Advanced Function Version 2 Release 1 and VSE/SP Version 2 and Version 3
- Transaction Processing Facility Version 2 Release 3.

The PR/SM feature allows new uses of 3090-E systems. Installations requiring multiple images now have four choices. There are four options for an installation with a workload requiring two independent systems each about the size of a 3090-200E. They are:

- Two separate processors such as two 3090-200Es
- A partitionable processor such as a 3090-400E
- A 3090 PR/SM processor such as a 3090-400E
- Use of VM on a processor such as a 3090-400E.

The decision must be made based on the trade-offs between isolation and versatility (that is, buying hardware rather than using software for device sharing).

The following sections describe PR/SM as it is known today, the new event-driven PR/SM dispatcher, how PR/SM evolved from 3090 multiple high performance guests support feature and performance considerations and application strategies.

Description of PR/SM

Central Processors

Central processors can be dedicated to a single logical partition or shared between multiple partitions. For example, on a 3090-400E, CP1 can be dedicated to LPAR 1 while CPs 2, 3 and 4 are shared by LPARs 2, 3 and 4. The exact details of implementation and terminology are unknown at this time. Operator commands at the 3090 system console control the sharing of the CPs between the LPARs with weighting factors and priorities. The weighting factors are used by the PR/SM dispatcher to compute target CP resource consumption. These are targets for CP consumption because PR/SM will allow an LPAR to use more CP resource than its target.

See 3090 PR/SM *page 38*

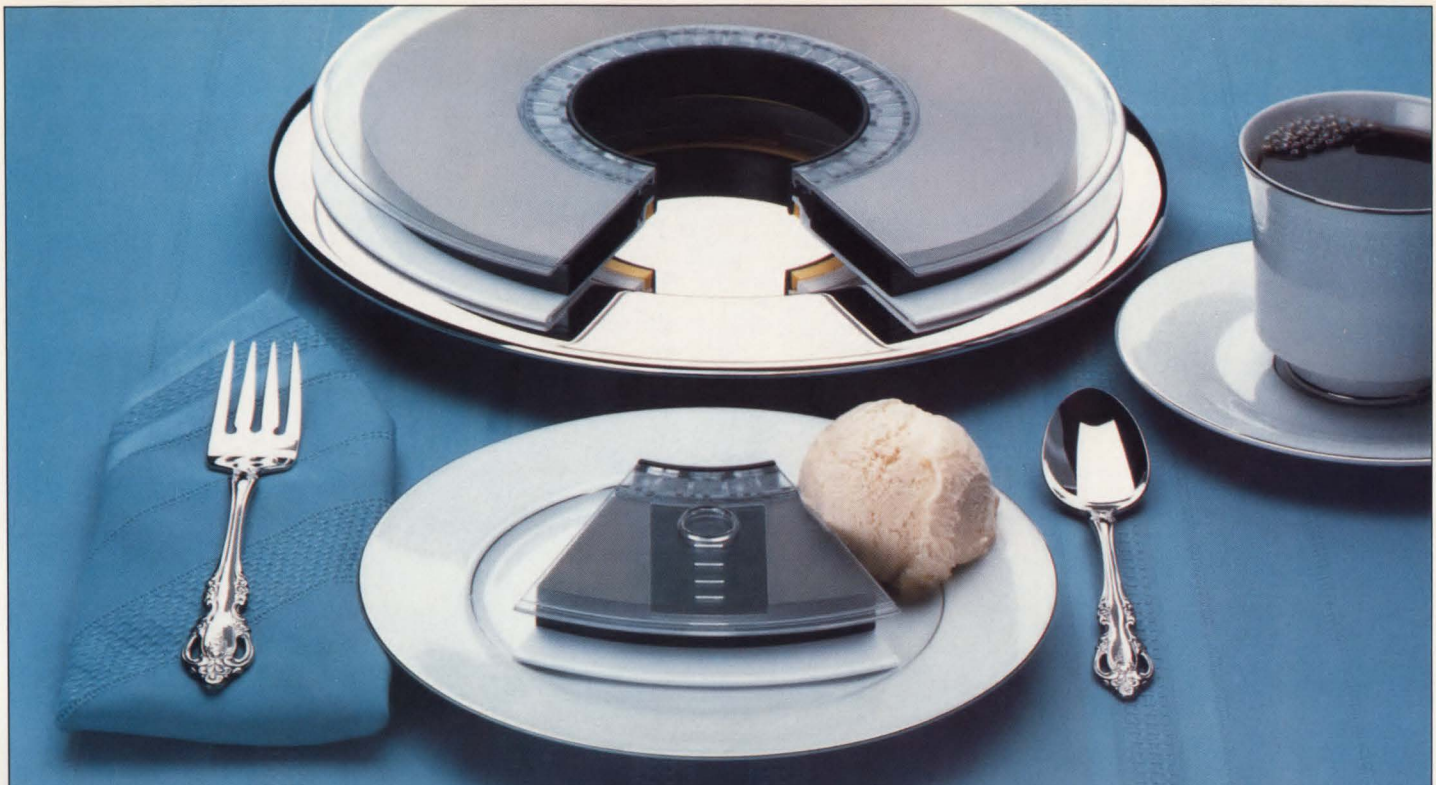

When choosing a network access package...why settle for just a piece of the pie?

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

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

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

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

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

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

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

Westinghouse
SOFTWARE **SOLUTIONS**

Westinghouse

CIRCLE #128 on Reader Service Card ▲

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

SIMPLIFY.
PROTECT.
CONTROL.

Introducing Multi-image Manager. To simplify, protect, and control data integrity, tape drive allocation, and console operation in the growing complexity of today's multiple image or CPU environments.

Multi-image Manager

Formerly SIS/SDM *Products*

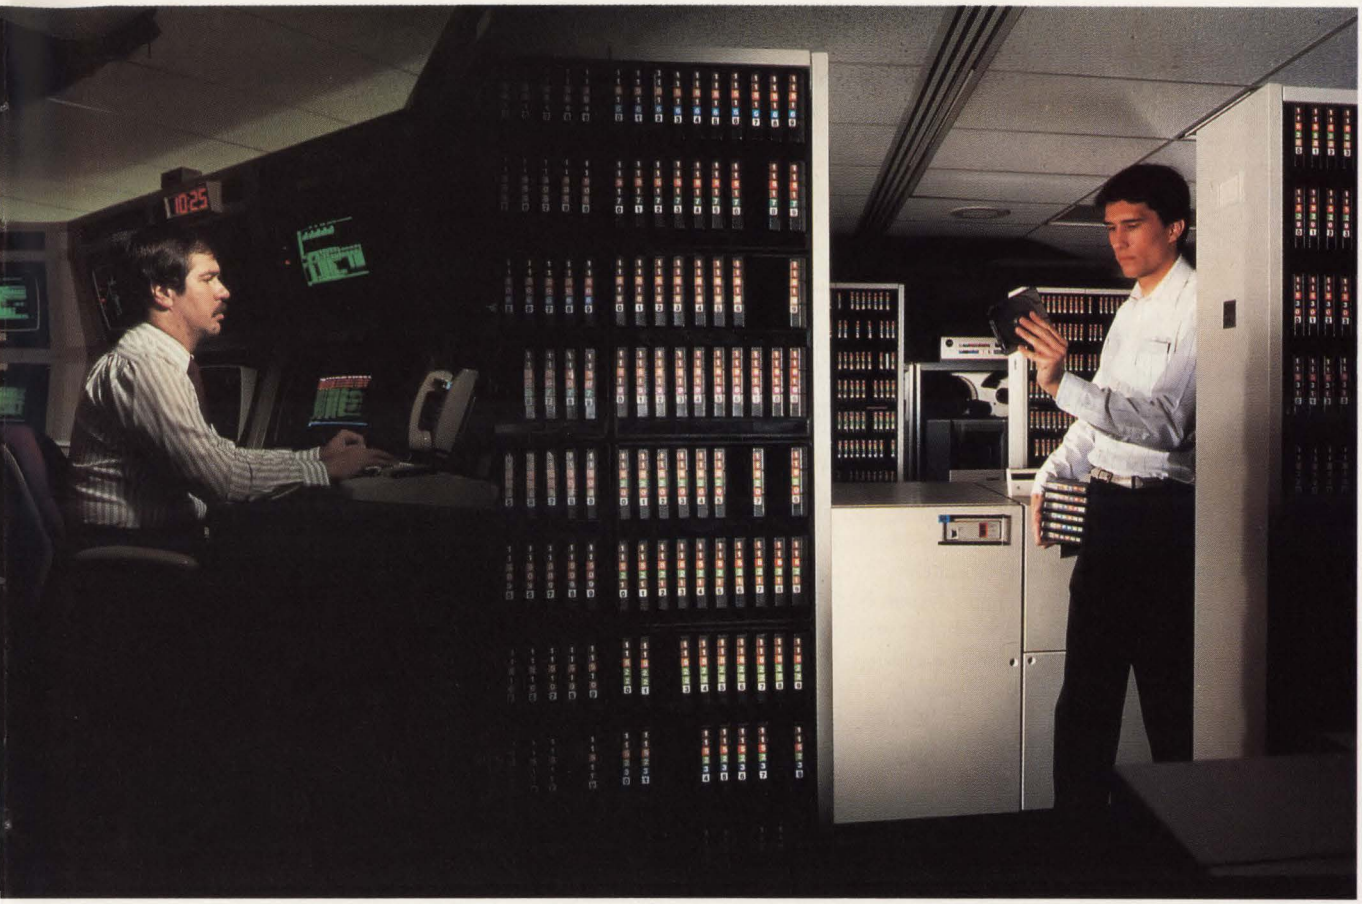

The efficient, cost-effective way to operate the multiple MVS environments of today is by sharing DASD, tape devices, and consoles. In fact, a majority of all data centers share resources to unify operations, provide backup, and defer hardware costs.

But sharing resources without protection or control dramatically increases your chances for deadlock conditions and data corruption. And physically separating or manually tending to your DASD and tape devices can cut deeply into your shop's productivity, and your budget as well.

Thousands of financial, manufacturing, retail, government and non-profit organizations, large and small, have decided they want increased productivity without risking the loss of data integrity. They operate their data centers with shared resource management products from Duquesne Systems.

Now there's a better reason than ever to eliminate risk and increase productivity in your data center. **Multi-image Manager** combines the performance of SIS* and the functionality of SDM into one new superset product.

Multi-image Manager offers all the benefits you need for a data center that runs more smoothly, efficiently, and with greater reliability. And that means greater peace of mind for you.

Here's how Multi-image Manager simplifies, protects and controls your operation:

- By providing data set integrity and device allocation across all system images in the complex. *This means fewer reruns to repair damaged data sets, and increased availability of* DASD *and tape drives.*
- By reducing complexity and increasing productivity for operations personnel. *Operating a multi-image/*CPU *environment as if it were a single system means fewer decisions, fewer errors and increased throughput.*
- By offering the opportunity to reduce costs by eliminating redundant hardware. *Safe sharing of resources between images or* CPUs *means that you can operate more effectively with less hardware.*

Multi-image Manager installs quickly without the need to IPL, incurs little overhead, and quietly guards against accidental destruction without interruption.

More than 1,200 Duquesne Systems customers successfully use SIS and SDM. Now *you* can enjoy the combined benefits of SIS and SDM in Multi-image Manager to simplify, control and protect your data center to its fullest. Call your local Duquesne Systems office, or call toll-free **800-323-2600 (in Pennsylvania 412-323-2600).**

*SIS (Single Image Software) and SDM (Shared Device Management) are two Duquesne Systems products combined and enhanced to form one new superset product, Multi-image Manager.

DUQUESNE SYSTEMS
Two Allegheny Center
Pittsburgh, PA 15212
Telephone 412-323-2600

CIRCLE #1 on Reader Service Card ▲

3090 PR/SM *from page 34*

PR/SM has been designed with the following objective: to maximize throughput by allocating all available physical CP cycles. If a logical partition does not use all of its allocated CP resource, PR/SM will allow other logical partitions that are sharing CP resources to use the excess. Operator commands can be used to change the CP target percentages while the logical partitions are active.

If an optional Vector Facility (VF) has been installed on a CP, then the VF is available to all partitions that will execute on that CP. When a CP is dedicated to a logical partition, its associated VF is only available to that logical partition.

Channel Paths

Channel paths are individually defined to a logical partition. They are not shared between logical partitions, but channels can be dynamically reconfigured between logical partitions. In this way the channels are directly controllable by the operating system using a logical partition. A device can be shared between logical partitions by using a separate channel path from each logical partition.

The implementation of the channels is one of the main differences between VM/XA and PR/SM. VM/XA allows sharing of devices and sharing of channels. VM/XA must then simulate some of the channel related I/O. PR/SM isolates the channels so that the operating system within each partition owns the channels directly.

Central and Expanded Storage

Central and expanded storage are defined to logical partitions prior to partition activation. When a logical partition is activated, the storage resources are allocated in 1M contiguous blocks. Memory resources are not sharable between logical partitions. Modification of the storage partitioning is possible, but only by deactivating and reactivating some logical partitions. As each partition is activated, storage is divided up sequentially.

Comparison to Amdahl's Multiple Domain Facility

Although significant efforts have been made to differentiate PR/SM and Amdahl's Multiple Domain Facility (MDF) in the marketing information that IBM has provided, these two offerings must currently be viewed as equivalent until actual measurement data is available for the IBM implementation. However, when comparative benchmarks are conducted, it will be interesting to evaluate the relative overheads of the two vendor alternatives.

Based on the available technical descriptions of PR/SM, one unique feature appears to be the availability of event-driven scheduling of the processors. Currently, MDF offers only a time-sliced implementation that Amdahl is currently describing as superior. History has shown that Amdahl's macro code implementation and desire to maintain its customer's residual value should allow them to provide a similar implementation if the facility addresses customer needs.

Event-Driven PR/SM Dispatcher

The dispatcher was designed with two objectives: first, to maximize throughput by allocating all available physical CP cycles and second, to maintain I/O responsiveness.

The dispatcher has many separately dispatchable units. Each logical processor of each partition-sharing CP resource is considered a separate dispatchable unit. Even though only two CPs exist in a 3090-200E, if four logical partitions are activated all sharing the two CPs, then there are eight dispatchable units. Logical processors from different logical partitions can be dispatched concurrently. It also means that all processors assigned to a single partition may not be active concurrently.

The dispatcher uses a dynamically-determined dispatch interval or time slice interval to establish the maximum time a logical processor can remain active for a single dispatch. The user can set a weighting factor from the systems console at activation for each partition which establishes a performance policy used by the dispatcher in selecting the dispatch sequence for logical processors. The weighting factors can be modified from the systems console while the logical partitions are active. Several classical techniques are used for dispatching and scheduling for the dispatcher to achieve its objectives.

When a logical processor enters an enabled wait state, the dispatcher will select another logical processor to run. This technique helps maximize throughput by maximizing the use of available physical processor cycles and improves I/O response times.

At the end of an interval, an active logical processor will be pre-empted and the highest priority ready logical processor will be dispatched.

When an I/O interrupt occurs for a higher priority logical processor, the lower priority logical processor is pre-empted and the higher priority logical processor is dispatched. This technique helps improve I/O response times.

Performance Analysis

Specific performance information regarding PR/SM will not be available from IBM until the third quarter of 1988. The performance of logical partitions will depend on the number of partitions, their configurations, the operating systems and workloads, the base machine configuration and performance tuning parameters. For interactive workloads such as IMS/VS and TSO under MVS/XA or CMS under VM/SP HPO, IBM is projecting the aggregate performance to be in the range of 88 to 110 percent of a single system operating in S/370 or ESA/370 mode on the total configuration.

When large single-image processing is not required, increased throughput is possible due to the reduced functional requirement.

IBM's prototype measurements of a laboratory MVS/SP 2.2 IMS workload on a 3090-200E configured as two equal dedicated uniprocessor logical partitions achieved an ITR that is 103 percent of that of the same workload operating in 370-XA mode on the entire dyadic.

When logical partitions share CPs rather than have dedicated CPs, more utilization will be required for PR/SM dispatching. However, the event-driven scheduler of PR/SM should recognize unused CP resources and balance the workloads resulting in an improved throughput. The use of the PR/SM feature with shared CP resources may improve responsiveness and throughput especially in an environment with workloads that exhibit fluctuations in processing demands.

Applications Strategies

Communication Management Configuration Host

For installations considering implementation of a Communication Management Configuration (CMC) host, PR/SM provides an interesting option. Previously two straightforward options existed. First, buy another processor with the required software. Second, VM could be used. Most installations that have this implemented today run the CMC on a separate hardware and software system largely due to the greater isolation provided by separate machines. The alternative of using

physical partitioning is generally not attractive because the CP resources required by a CMC are typically in a small percentage of an installation's workload.

PR/SM would allow a logical partition to be defined as the CMC host. This does not offer the isolation of a separate system, but could be very cost effective. A second system requires a second license of the SCP, VTAM, JES, RMF and so on. The CMC host could be defined with more CP resource than it normally requires. At network startup, additional CP is required and it will be available. After the network is initialized, the PR/SM dispatcher will allow other partitions to use the CP resource not required by the CMC host during normal operations.

Partition Memory Asymmetrically

Installations with physically-partitionable processors have always had the dilemma of one side needing more memory or channels than the other. Physically-partitioned processors have to be fully symmetrical in respect to channels, memory and CPs (except the new 3090-500E that has three CPs on the side zero and two CPs on side one). In the past, an installation needed to add resources to both sides of a processor to meet the requirements for symmetry.

Consider an IBM 3090-400E with 128M central storage and 128M expanded storage running in physical-partitioned mode with a large TSO workload on one side and a large CICS workload on the other.

TSO requires expanded storage and CICS performance is enhanced by adding central storage. Which should be planned for? Installing PR/SM in this environment adds versatility since resources can be split unevenly, while maintaining dedicated CPs for similar throughput. The cost of PR/SM is less than the cost of either central or expanded storage for this installation.

PR/SM and the Extended Recovery Facility

The Extended Recovery Facility (XRF) provided by MVS/XA, VTAM, NCP, IMS and soon by CICS/MVS is best served by two physically separate systems. The following is a typical XRF implementation scenario.

For discussion purposes, consider two separate systems running the appropriate levels of the XRF software. They can be called System A and System'B in which System A will be the primary IMS/VS on-line system. System B is the backup system normally running a TSO development workload. IMS/VS is also running on System B monitoring System A to determine if an XRF takeover is required. If an XRF takeover occurs, then MVS/XA on System B will give high priority to VTAM and IMS on System B to provide enhanced performance during the takeover.

The disadvantage of this approach is that System B, the backup system, will often be running XRF software with different maintenance levels than System A. This is because many installations use their development system as a test system for their system software. When an XRF takeover migrates the System A workload to System B, the actual software may be different. Sometimes this is planned so that if the new software has problems backing out, the changes can be done with an operator console-initiated XRF takeover to move the workloads to the previous systems software environment.

With PR/SM the System B could be a logical partition running the same software as System A. System C, another logical partition sharing CPs with System B, would be used for the TSO development system. Normally when System B is monitoring System A, its CP requirements are small leaving the CP resource available for System C's TSO development workload.

When a takeover occurs, System B will suddenly have a large CP resource demand. If System B has been activated with appropriate weighting factors and priorities, then the PR/SM scheduler will take the CP resource from System C for System B to use during the takeover. In this way System C, the testing system, is clearly isolated from the main System A IMS on-line and System B, the XRF backup.

Impact on Software Licensing Fees

Although PR/SM is a logical response to customer requirements and market share earned by Amdahl's MDF offering, PR/SM has one interesting consequence for IBM. At a time when IBM's software investment has never been higher (that is, ESA, DFP, SAA, OS/2 ...), IBM has provided the large user with a way to reduce long-term license fees. As such, it would not be surprising if IBM chose to increase base license fees for processors with MDF or PR/SM-like facilities.

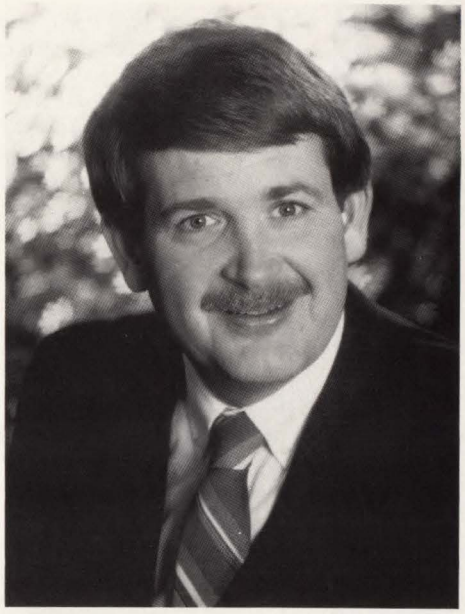

H. Pat Artis

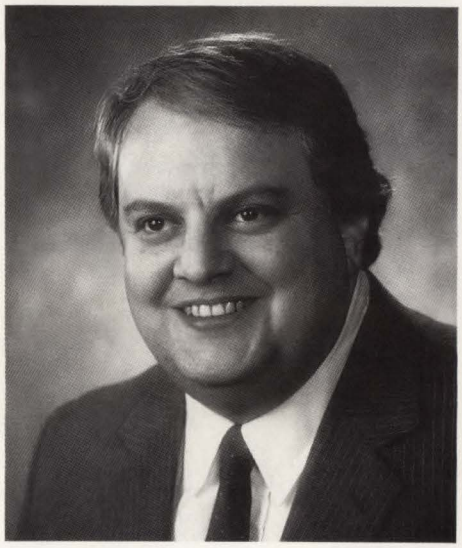

Alan M. Sherkow

ABOUT THE AUTHORS

H. Pat Artis is a recognized authority in workload characterization, forecasting simulation modeling and capacity planning. He is the author of more than 75 papers and lectures internationally. Artis is president of a private consulting firm, Performance Associates, Inc., Vienna, VA 22180, (703) 281-7399.

Alan M. Sherkow is president of I/S Management Strategies that specializes in strategic and tactical planning. Previously, Sherkow was the lead capacity and performance planner at the First Wisconsin National Bank of Milwaukee, WI. He received his BS and MS degrees in computer science from the University of Wisconsin, Milwaukee.

TECH ADVISOR
Answers From The

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 *We are an IBM 4381 shop running VSE 2.1, VTAM and CICS shutdown statistics entry titled, "Times Storage Recovery Entered." This number is consistently around 3,000 while the "No Storage Violations" message is displayed below. I understand that CICS believes it recovered any broken storage chains, but I feel that there is some corrupted storage left from the recovery. How can I identify the transaction(s) that is causing this? I should add that the total number of storage acquisitions/releases is averaging 5.5 million and total tasks average 220,000 per day.*

A Judging from the information that you have provided, your problem sounds like the Dynamic Storage Area (DSA) rather than a particular transaction. The System Initialization Table (SIT) contains a value for the Storage Cushion Size (SCS). The SCS is used to prevent a Short On Storage (SOS) condition. When the amount of free storage in the DSA falls below the storage cushion, CICS will release all non-resident programs that are not in use. This process could cause the situation you have described.

You will need to either increase the size of the DSA or decrease the size of the SCS and see the CICS installation and operations guide for instructions on how to determine the correct SCS and the correct DSA size for your installation.

If you are running CICS under ICCF, then the DSA size is controlled by the size of the ICCF partition "T0." In this case to increase your DSA size, you would increase the T0 partition size (which is an ICCF partition, not a VSE partition). If you are running CICS stand-alone, then the DSA size equals the DFHSIP execution size minus the CICS nucleus size. For example: size = 8192K, nucleus = 500K, DSA = 8192-500.

Although the problem you describe should not cause any storage corruption, it does result in performance degradation. Adjusting your DSA should reduce the "Times Storage Recovery Entered" number and in addition give you an added performance benefit.

Q *We are running two 4341-12s, DOS/VSE 2.1, non-VM machines using IBM's Shared Power on our DASD spool file. We are experiencing a high level of I/O waits due to the contention of the lock file because of this share-protect feature. Please advise if there is any known solution to this problem.*

A Your question regarding the VSE lock file is one that is very common to many VSE shops.

First, make sure you have covered the simple solutions. The two most highly used files in your system will be the lock file and the VTOCs. So the first thing to do is to put the lock file on the least-used DASD, even on a device by itself. The next questions are: do you have to share POWER, what is this worth to your shop and can you get by with separate POWERs. Given that the answers are no, then we look at the third solution.

There are software products on the market to deal with this issue such as CACHE/MASTER from SDI or Extend/VSE from Goal Systems. What these products do is to put the lock-file in memory and therefore reduce real I/O time. Unfortunately they require VM, so this solution is out.

The last possibility is to review POWER's PNET. To use this you will need a 37XX communication controller and channel attach both CPUs and use PNET to route print output to the CPU that has the printer. This solution works if you are sharing only POWER and not other files. Of course, POWER is the main contender for the lock-file and removing it from the picture would be quite a bit of relief.

Q *Subsystems — What are they? Why are they so powerful? Why are they used by RPT, ACF2, TMS, etc.? Regarding disaster recovery — how do people keep catalogues and the TMS, TMC backups reasonably current?*

A There are many approaches to handling disaster recovery backups. However, the steps below should provide you with a reasonable level of protection while allowing you to easily restore either elective datasets or entire packs rapidly and easily. These procedures assume that you have some sort of data management utility in-house (such as DF/DSS, etc.), but you can use the basic philosophy even if you do not.

1. Take a nightly backup of the TMC, using the TMS TMSCOPY utility. This will back up the TMC in the proper format and will also reset the TMS Audit dataset. You should retain at least eight generations of this backup.

2. Take a "cumulative incremental" backup on a nightly basis. A cumulative incremental is a backup of all datasets on the system which have been changed since the last full-pack backup. This can be accomplished by backing up all datasets with a DSCHA flag of one.

DO NOT reset the DSCHA flag to zero on the nightly incremental backups. This will ensure that each successive nightly backup will include all datasets changed since the last set of full-pack backups were taken, making selective dataset restoration a straightforward, one-step procedure.

These backups should also be kept for at least eight generations.

3. On a weekly basis, take a dual-image, full-pack backup of all packs on the system. Reset all of the DSCHA flags for each pack to zero as the backups are completed, thereby preparing for the next week's set of cumulative incrementals.

The reasons for taking dual-image backups is twofold:

1. One copy may then be sent out for offsite storage while its counterpart remains onsite for quick access in the event it is needed for a restore. This will also save in offsite "emergency delivery" fees.

2. A backup copy is available in the unfortunate event of tape media failure during a restore.

Retention of the full-pack backups is up to you, but we always like to err on the safe side. It all depends upon your business and archiving requirements. I suggest at least two months worth and maybe more depending upon your worst-case scenario.

One other note: if you are using the data compression feature of DF/DSS (and you should, if possible), DO NOT use compress on the backups of your SYSRES volumes. This feature will make them unusable for a standalone restore, if needed.

Q *We run a 4341-M2 with VM/VSE/SP on 3370s. In order to improve the performance of the machine, we want to install 3380s. The question is which model of the 3380 controller do we install? AA4 or A04? (We have a 3880-3). The AA4 with Dynamic Path Selection and two controllers seems obvious, but Dynamic Path Selection and Dynamic Path Reconnection are MVS only. It appears that the A04 controller is all that VSE/SP can make use of. Help?*

A We agree that you will get better performance as well as increased capacity from 3380s as opposed to 3370s. The AA4 controller vs. the A04 model is a better choice because it provides you with two paths to your DASD. But because of the unit's limitations, you are forced to keep your data strings short to prevent bottlenecks. To address this problem, IBM has brought out newer controllers such as the AE4, AJ4 and AK4.

As far as the AA4 you asked about, it will work in your VM/VSE/SP environment. It will also work with Dynamic Path Selection and Reconnection if you ever migrate to VM/XA. IBM even maintains that the Dynamic Path features work *better* with VM/XA than they do with MVS/XA.

So, my recommendation is to install the AA4 controller rather than the A04, since it is about the same price as the A04 and it will perform better.

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

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

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

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

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

LRS

Levi, Ray & Shoup, Inc.

Specializing in Computer Systems Software

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

Telex (510) 60-00675

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

CIRCLE #121 on Reader Service Card ▲

Unlock The Mysteries Of DB2 Performance.

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

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

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

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

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

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

!Candle®

CIRCLE #120 on Reader Service Card ▲

DB2 Version 2

Offers Major Functional and Performance Enhancements

By Joel Goldstein

The next release of DB2 has an announced availability date of October, 1988 and offers significant extensions of the relational function, as well as operational and performance enhancements.

Relational Enhancements

Referential integrity is a *major* functional enhancement of DB2. It basically completes the required features for a relational database system and will improve the productivity of application developers. Referential integrity maintains the relationship of data values between related columns of different tables.

Using the tables in Figure 1 as an example: the Employee Table contains the column DPT_NMBR for each employee and the Department Table contains the column MGR_NMBR that is the employee number of the department manager. Make the obvious assumption there must be a matching DPT_NMBR in the Department Table for each unique entry in the corresponding column of the Employee Table. On the same basis, there should be a matching EE_NMBR in the Employee Table for each MGR_NMBR in the Department Table.

These are referential constraints. The Department Table uses department number as the primary key and manager number as the foreign key (manager number is an employee number). When looking at the Employee Table, the employee number is the primary key and the department number is the foreign key.

Referential constraints are enforced by DB2 during LOAD, UPDATE, INSERT and DELETE functions. The DB2 system will ensure that all primary key values are unique and that all foreign keys *always* have a corresponding primary key. One of three specific delete rules may be specified when the relationship is defined:

- SET NULL will allow the deletion of a primary key value and set all dependent foreign keys to null
- CASCADE will allow the deletion of a primary key value and will delete all dependent foreign key values
- RESTRICT will not allow the deletion of a primary key that has dependent foreign keys. *This is the default.*

Normally, for this standard type of example, two basic rules should be enforced: all employees should have valid departments and departments with employees cannot be deleted.

However, we all know that occasionally departments may not have a manager. If the employee who is now the manager leaves the company and is deleted from the Employee Table, the manager column for that department must be set to nulls. In the same manner, when a manager is appointed and his employee number is added to the Department Table, the system will verify that the employee number exists in the Employee Table.

Prior to the support of referential integrity, these relationships were maintained by the application that added to the complexity of function and code. Now the system maintains the referential integrity and the application development task is simplified.

One consideration of additional functionality is often overlooked. As systems provide *more* functionality, they grow larger and often consume more resources. Since DB2 Version 2 will not be available for several months, we can only guess if referential integrity will consume more resources. Hopefully, the elimination of application code and multiple SQL calls will more than offset the cost of using this new capability. It is also quite likely there may not be an easy way to measure the cost of this function since the overall throughput capability of DB2 has shown significant improvement.

Performance Enhancements

The benchmark numbers released by IBM show substantial improvements for all areas of DB2 processing. The performance improvements are due to improvements in several areas:

- Decreasing the instruction path length for create thread, sign on and authorization checking
- Improved locking methodology for the IRLM
- Improved sorting algorithms

FIGURE 1

DPT_NMBR	DPT_NAME	MGR_NMBR
ACCTG01	ACTG/PAY	149631234
ACCTG02	ACTG/RCV	136623456
.....	.....	.....
PURCH01	PURCHASE	151221234
SALES01	SALES	139698779
SHPNG01	SHIPPING	148635678

Department Table

Primary Key-DPT_NMBR
Foreign Key- MGR_NMBR

DPT_NMBR	NAME_LAST	PAY_GRADE	EE_NMBR
ACCTG01	BERMAN	22	149631234
PURCH01	CAMPBELL	16	139631734
.....	.....	..	.....
ACCTG02	FERGUSON	22	136623456
.....	.....	..	.....
PURCH01	JAMES	22	151221234
SHPNG01	KLONES	22	148635678
ACCTG01	LEPPER	18	143635679
SHPNG01	PAULSON	19	137635679
ACCTG02	RAMIREZ	19	147635670
SHPNG01	STRUK	20	148633456
.....	.....	..	.....

Employee Table

Primary Key-EE_NMBR
Foreign Key-DPT_NMBR

Referential integrity maintains the relationship of data values between related columns of different tables.

- Usage of sequential prefetch for temporary work files
- Elimination of logging requirement for temporary work files
- Improvements to the Optimizer that provide improved access path selection and reduced usage of workfiles
- Addition of a new index to SYSCOLUMNS
- Catalog statistics for non-indexed columns and a *cluster-radio* column in SYSINDEXES that aids the Optimizer during access path selection
- The ability to update statistics columns in the catalog which influence the access paths chosen by the optimizer.

The transaction throughput capability has improved to 438 per second for IBM's standard Credit Authorization Checking benchmark. This was run on a 3090-600E with 256k Real and 256k Expanded Storage using IMS/VS FASTPATH and MVS/ESA. A more standard transaction processing workload using IMS/VS full function as the Data Communications front end was able to reach a rate of 186 transactions per second compared to a rate of 123 transactions per second using the prior release of DB2 on a MVS/XA system — more than a 50 percent improvement to overall throughput.

The improvements to the Optimizer and sorting algorithms have shown that the consumed CPU time is reduced by more than 50 percent and execution times have been up to *four* times faster for some queries involving large sorts. Additional enhancements to data movement facilities and improved usage of cross memory services further reduces the CPU time for retrieving many columns of data from large tables. This will especially benefit batch reporting jobstreams by significantly reducing their CPU resource consumption and elapsed time requirements. The additional information available from the new Catalog columns allows the Optimizer to select better access paths and eliminate unwanted rows earlier in the retrieval process. The Optimizer now closes Transitive Predicates when creating the plan from an SQl request. A predicate like:

. . . where A = B and B = C

will now be properly closed by the Optimizer to reflect the obvious

. . . where A = B and B = C and A = C

This closure of the predicate provides improved evaluation and selection processing and can eliminate merging and sorting of work files.

The Create Index operation provides dramatic performance improvements. IBM has reported that the creation of a four-column index on a table of 255,000 rows consumed 35 percent less CPU time and executed *nine* times faster on a 3090-300E. This will certainly improve the ability to perform maintenance on large tables within limited operational windows.

The CICS Attach function is enhanced to allow an application to dynamically select a plan for execution. This enables an application to use smaller plans instead of requiring all possible plans to be bound together. The smaller plans require less time to load, less space in the pool and facilitate maintenance of the application since only individual plans that are changed need to be rebound.

Operational Enhancements

DB2 now provides a governor that is called the Resource Limit Facility (RLF). This allows an installation to exercise control over the CPU resources consumed by a dynamic SQL statement. The prior inability (in previous releases) to control the resources that might be consumed by "runaway" dynamic SQL requests was a major obstacle preventing installations giving a dymamic SQL capability to end users. The RLF provides a wide flexibility for establishing the criteria for determining resource limits. The established limits may be changed dynamically while the DB2 system is running.

Tablespaces may now be segmented. This provides improved operational, performance and space management facili-

ties and makes it easier to manage the environment. A tablespace can be divided into groups of pages called segments. A segment may be specified for specific tables as needed. Unlike previous releases of DB2, the rows of a table will occupy only the pages of its segment and will not be intermixed with rows of other tables. This allows for locking to occur at the *table* level rather than the tablespace level. This will provide improved concurrency when multiple tables exist within a single tablespace. The individual tables may be reorganized without affecting the organization of the other tables within the tablespace. Space from dropped tables may be reclaimed without having to reorganize the entire tablespace.

An optional AUDIT facility provides an audit trail of specified DB2 events. The facility logs specifically requested activities that can include security violations, access to sensitive data and the results of GRANT or REVOKE authority requests.

Improved Recovery extensions now provide for point-in-time consistency of one or more tablespaces and provide reports containing the necessary data for recovering one or more tablespaces. This is a required complement to referential integrity that allows the recovery of multiple tables to a consistent point-in-time.

Summary

This information has been summarized from several public presentations by IBM personnel and other formally released documentation. There are some additional new features and support enhancements that are not covered here solely because of my personal viewpoint and selection of the areas I thought were of greatest significance.

DB2 Version 2 is a major step forward both in function and performance.

ABOUT THE AUTHOR

Joel Goldstein is president of The Responsive System Company, an information systems consulting firm that specializes in on-line database systems. He has 25 years of data processing experience and has been deeply involved in the design and performance of database systems for 18 years. He has held management positions at major communications and insurance companies and has presented papers on IMS Performance Tuning at SHARE. Responsive Systems Company, 81 Grand St., Jersey City, NJ 07303, (201) 333-7272.

Software Review:

Monitoring VM Performance

Just the Ticket for Vital Signs

By John Kador

Minnesota's largest city has seen more ups and downs than the Dow Jones Average. Home to such major employers as Control Data Corporation (CDC) and Cray Research, Minneapolis is an ideal breeding ground for innovative software. Thus, when a 1985 shutdown of CDC's IBM plug-compatible disk operations offered employees new opportunities, it made possible the creation of BlueLine Software and the marketing of its VM performance monitor, Vital Signs, the subject of this review.

Vital Signs was announced in July, 1987. BlueLine acquired the rights to three performance monitoring software packages from Richard Jensen, an independent consultant and systems developer and former member of the IBM team that developed the VM operating system. Jensen also played a significant role in the development of IBM's VMMap and Smart performance monitors.

In announcing Vital Signs, BlueLine hardly staked out virgin territory. Besides the IBM products mentioned earlier, a number of independent vendors offer VM performance monitors including Explore/VM from Goal Systems (Columbus, OH) and Omegamon from Candle Corporation (Los Angeles, CA), among others.

Output Displayed Graphically

Like VMMap and Smart, Vital Signs collects statistics about key systems performance variables such as CPU utilization, resource waits, I/O and paging rates. But unlike its IBM counterparts that present their numbers in a tabular fashion, Vital Signs displays its output graphically.

For Memorex Corporation's Communications Group in Milpitas, CA, the graphs simplified the interpretation of its 4381's raw performance data. Without such data, the CPU was able to operate at only 20 percent of its potential capacity according to a Memorex systems engineer. Today, that same processor is running at 60 to 80 percent of maximum efficiency because, as the systems engineer recalls, "Most of our system peaks were occurring on the same disk controller and channel." In response, the company simply added more devices, channels and controllers to its configuration to ensure that the user load is spread more evenly across the system. Memorex is convinced that the graphical output provided by Vital Signs made the interpretation of the problem much easier. By contrast, the burden of analyzing raw statistics presented in tabular format would have retarded the diagnosis process.

Real-time and Historical Reporting

Vital Signs provides both real-time and

"With VM FAX, I can turn any VM terminal into a fax machine!"

Data processing managers and PROFS users everywhere are enjoying the power of VM FAX—the first communications software package to let you send faxes from your IBM mainframe and mid-range terminals.

And, if you have a graphics terminal, you can send and receive graphic images as well as alphanumeric fax messages. In fact, many businesses are considering replacing all or most of their machines with one graphics terminal, printer, and operator for greater effectiveness.

VM FAX is so easy to use. Messages can be sent using the PROFS NOTE menu. When sending a message, just indicate your recipient by "nickname"—no long string of fax phone numbers to remember.

You'll also appreciate features such as automatic dial and re-try. When receiving messages, you can print them out at your printer or route them to another printer. With VM FAX, transmissions can even be stored and retransmitted later. Soon you'll even be able to display them on an IBM 3193 "all points addressable" terminal.

VM FAX is really a complete message management system. As part of the VM MESSENGER family of software products, it allows DP managers and other users to combine electronic mail, subscriber databases, telex, and packet switching into one cohesive system—with a single user interface and optional integration with IBM's PROFS office automation product.

Best of all, it's another highly-acclaimed product from Systems & Telecoms—a leading international developer of communications software for the VM/CMS and UNIX operating systems.

For a FREE, 30-day evaluation copy of VM FAX, write or call us today.

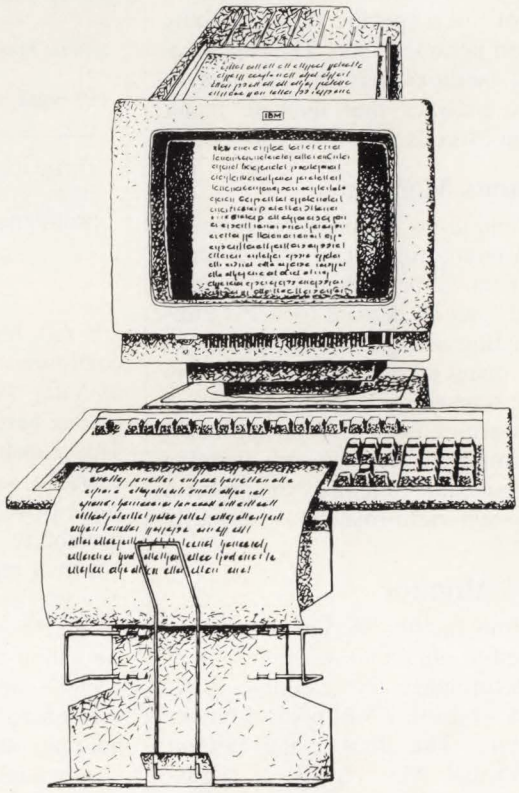

SYSTEMS & TELECOMS

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 ▲

historical reporting measuring workloads and performance in the VM operating environment to enable the user to eliminate bottlenecks and plan for future equipment needs. The tool can help MIS managers and systems programmers in optimizing current performance and in getting a better handle on capacity planning. The useful life of hardware can be extended not only by identifying and eliminating the bottlenecks, but also by predicting the effects of changing workloads and various equipment alternatives using Vital Signs' modeling features.

Vital Signs is comprised of a Continuous Monitor, Intensive Monitor and Historical Reporting facilities. The Continuous Monitor is used for long-term data gathering and real-time access to displays of all major performance variables. The Intensive Monitor provides both a DASD modeling facility and the ability to take a "snap shot" of a specific system activity over a short period of time. The Historical Reporting facilities perform and analyze long-term analyses that indicate major performance trends within the data center.

Continuous Monitor

The Continuous Monitor provides an exception monitoring facility that notifies users of potential problems as they occur. Thresholds can be defined for CPU utilization, paging rates, free storage extend rates and many others. When the system exceeds a threshold, Vital Signs sends a message warning the user of an approaching performance issue. The early warning gives the user time to remedy the problem before system performance is seriously degraded.

Intensive Monitor

A real-time facility, the Intensive Monitor, is used to obtain an in-depth look at system performance. For example, it allows users to display VM control blocks in real time. The three major control blocks (Virtual Machine Block (VMBLOK), Virtual Device Block (VDEVBLOK) and Prefix Storage Area (PSA)) can be decoded and displayed (see Figure 1). This ability aids in problem determination. For a given virtual machine, the VMBLOK contains such information as the dispatch and priority level of the virtual machine, the virtual machine's processor registers and the options currently in effect. The VDEVBLOK provides I/O counts by virtual device. The PSA is the root of all control blocks in the system.

Unlike other performance monitors that are too resource-consumptive to allow continuous monitoring, the low overhead of Vital Signs permits users to monitor system performance on an ongoing basis. This capability eliminates the need for users to sample performance data over short periods of time and minimizes the possibility that the periods sampled are not truly representative of the data center as a whole. It also makes possible instant access to both current and historical information based on any point in the data center's operations. Finally, continuous monitoring makes exception monitoring possible, as well.

A modeling feature useful for "what-if?" analyses is included in the performance monitor (see Figure 2). The modeling function can be used to test the effect of equipment changes on actual workload without the user having to install equipment to re-enter the data describing the current job stream. In this way, "what-if?" analyses can be performed for different storage devices, channel loads, file block sizes and I/O loads. This capability aids MIS managers because the impact of a spectrum of alternatives can be compared in minutes to provide detailed performance implications of both current and future equipment strategies.

FIGURE 1

Display Control Blocks

```
06/16/87              I N T E N S I V E   M O N I T O R              VTLS0080
17:06:05                    VMBLOK FOR LLOYD                          MENU
------------------------------------------------------------------------------
Temporary Storage Size In Bytes: 13,369,344    Next VMBLOK in Queue: 003F5E20
Permanent Storage Size In Bytes:  3,145,728    Prev VMBLOK IN Queue: 003F60E0
Projected Working Set In Pages:          22    Next Cyclic User:     003F8590
Resource ID If 370X Terminal:          0000    ECBLOK Pointer/CREGO: 000000E0
Line End/Line Delete Characters:        #/¢    Segment Table ADDR:   0C3E59C0
Character Delete Character:               @    VCHBLOK Table ADDR:   003EE9F0
Accounting Information:           62770010     VCUBLOK Table ADDR:   003A1780
Current Distribution Code:        TECH/SNA     VDEVBLOK Table ADDR:  003E1348
Secondary Console Userid:                      Terminal RDEVBLOK:    00000000
  *** ATTACHED CHANNELS/CHSTRT INDEX ***       DISP to CON VDEVBLOK: 0000
  0000 0030 0060 0090 FFFF 00C0 FFFF FFFF      BASEADDR/Lock Holder: 00000000
  FFFF FFFF FFFF FFFF FFFF FFFF FFFF FFFF      ADDR EXT Trace Block: 00000000
Program Status Word:    FF060000 00165DB8      ADSTOP Control Block: 00000000

R0-7: 00000000 00004130 C3D6D5F1 001601BA 00000000 00000000 0000001C 00000000
R8-F: 0000026C 000022E0 00000010 000025F0 00160158 0000DFA0 7016041A 70165D1E
FO-N: 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
------------------------------------------------------------------------------
Press the FORWARD key to see the next page.                              024
=>
PF1=HELP PF2=EXPLAIN  PF3=RETURN          PF8-FORWARD PF11=SET PF12=PRINT
```

Three major control blocks — VMBLOK, VDEVBLOK and PSA, may be displayed. Here, VMBLOK for user Lloyd is decoded to make it easier to read and isolate problems.

Users Responses

One user, Overnight Transportation, based in Richmond, VA is the fifth largest trucking company in the United States. Another user is Valmont Industries, Inc., a diversified manufacturing company. Their experience with Vital Signs will shed some light on the product.

Overnight Transportation

This subsidiary of Union Pacific Corp. (NY) has more than 2,000 trucks serving 150-plus terminals in the U.S. and Canada. In 1987, the company rang up revenues of more than $800 million. At its Richmond data center, Overnight Transportation operates an IBM 3081 under VM running end-user applications such as IBM's PROFS. It also operates an IBM 3083 running VM and two guest operating systems. Third-party software includes V/Force and V/Seg both from VM Systems Group (Arlington, VA).

VM Systems Programmer Dick Davis joined Overnight Transportation in early 1987 and immediately wanted to find a

replacement for VMMap and Smart. After hearing about Vital Signs, he agreed to be a beta site. Subsequently, Overnight Transportation installed the product on a formal basis.

Its ease of use and resulting savings in manual labor was the main justification for Vital Signs. "It provided us with a way of looking at historical information as if it were real time," Davis says. He estimated that Vital Signs paid for itself by reducing manual effort well in advance of the 36 months used in the justification.

Overnight Transportation uses Vital Signs to assist with capacity planning questions, pulling data off during peak periods. That data is fed into a modeling program that IBM runs and is used to balance CPU loads. The flexibility of Vital Signs has been a boon to Overnight Transportation.

Unlike VMMap and Smart, Vital Signs allows users to cut and slice data in many ways. With the former products, it is often problematic to recapture data. "With VMMap and Smart, once we processed data, that was it — we could not go back. That is not a problem at all with Vital Signs," Davis notes.

While he cannot point to deficiencies in Vital Signs, Davis is quick to note that the product is still evolving. Some user screens that Davis desires, for example, are not available yet. These panels include page migration and alternate channel, both of which he gets from using Smart. Eventually, says a spokesman for BlueLine Software, Vital Signs will support the page migration feature, perhaps by the end of this year. As for alternate channel, there are no plans to include it.

When asked to note an interesting feature of Vital Signs, Davis mentions that on its DOS/VSE system, the company emulates 3340 storage devices on 3380 storage devices. "With Vital Signs, I am able to get such statistics as Start I/Os and Loads from emulated disks. With VMMap and Smart, it is not possible to get that information. It is not even possible to request it!"

Valmont Industries, Inc.

A diversified manufacturing company, Valmont Industries employs 3,500 people around the world — 1,000 of them in Valley, NE, a western suburb of Omaha. It operates an IBM 3090 180E running VM/SP and VM/HPO as well as five VSE guests and 50 CMS users.

In early 1987, Ned Hedrick, manager of Technical Services, was looking for a VM monitoring product that would provide two benefits. First, he wanted real-time monitoring of the VM system to take snapshots and determine execution characteristics. The system sometimes suffered from performance degradations presumably caused by bottlenecks and Hedrick wanted to eliminate them. Second, he wanted a facility to provide batch-oriented reports for management, primarily for capacity planning purposes. Valmont installed Vital Signs in November, 1987.

The company uses Vital Signs' bottleneck analysis more than anything else. According to Hedrick, "We like its real-time SEEKS analysis that allows us to look into the I/O subsystem to determine what is going on in a particular disk device." In this way, Vital Signs pointed out some real hot spots on Valmont's disk drives. "As a consequence, we moved datasets around and improved performance thanks to the visibility Vital Signs provided," Hedrick adds.

Has Hedrick seen any limitations of Vital Signs? "None!" Does he have a wish list? "Of course!" High on that list would be some automation of the performance monitoring process that would act to limit that particular user's impact on the system (if a particular user's virtual machine began to use too much of a certain resource). At this time, while Vital Signs can sound an alarm when a user-defined threshold is exceeded, the user must take appropriate remedial action on a manual basis.

Vital Signs integrates both real-time and historical reports presented in tabular and graphic display formats. The data extraction facility offers the flexibility of allowing data to be transported into external analytical software like SAS, Lotus 1-2-3 or Symphony. This capability allows users of Vital Signs to analyze summary results further and customize reports and graphs.

Its real-time displays and printed reports show all major performance measures including I/O rates, CPU utilization, resource waits and paging rates. The displays can be tailored to focus on specific time periods, users and other system components.

See Vital Signs *page 111*

F I G U R E 2

DASD Configuration Modeling

```
10/07/87                          V I T A L  S I G N S                    VTLS0003
11:50:38                    DASD CONFIGURATION MODELING                     ENTRY
------------------------------------------------------------------------------
Device to be modeled:   243                              Device Type:          3350
  ---- INPUT TO MODEL ----                           D  Average Seek Time:    .0250
Device Type:            3350                         A  Bytes Per Track:     19,069
Channel Utilization:     .05                         S  Transfer Rate:    1,226,752
I/O rate per second:       7                         D  Rotational Delay:     .0168
Average Block Size:    1,024                            Average Block Sizes: 1,024

<== C H A N G E S ==> <======== C A L C U L A T E D  R E S U L T S ========>

DEV  CH  I/O  AVG BLK    DATA      DEVICE    CONTENTION    EFECTIVE    RESPONSE
TYPE UT% RATE  SIZE   TRANSFERED SERVICE UT% RPS-M QUEUE SERVICE UT%    TIME
---- --- ---- ------- ---------- ------- --- ----- ----- ------- --- --------
3350 .05    4  1,024      4,096   .0342 .13 .0008 .0053   .0404 .16   0.1928
3350 .05    7  1,024      7,168   .0342 .23 .0008 .0107   .0458 .32   0.4716
     .00    0      0          0   .0000 .00 .0000 .0000   .0000 .00   0.0000
     .00    0      0          0   .0000 .00 .0000 .0000   .0000 .00   0.0000
     .00    0      0          0   .0000 .00 .0000 .0000   .0000 .00   0.0000
------------------------------------------------------------------------------
You may now change the model parameters.                                  096
=>
PF1=HELP PF2=EXPLAIN PF3=CANCEL                      PF11=SET PF  VITALS
```

Vital Signs' modeling facility permits "what-if?" analysis to explore the effect of various hardware and software alternatives on system performance.

Shared Segments in VM and MVS

By Ole Reitzel Jensen

There is a clear resemblance between DCSS and PLPA and they are both based on the use of a number of common page tables.

In VM related writings, the term Discontiguous Shared Segments (DCSS) is often used. The name refers to page segments outside the normal machine size, code segments used in common by a number of VM machines.

The primary purpose of DCSSs is to improve performance in systems in which many users use the same group of programs. With DCSS, only one copy of a given program needs to be in real storage and program load overhead is minimized.

In the MVS world, no one talks about DCSSs. Common code is placed in Pageable Link Pack Area (PLPA) where it can be referenced by all users (address spaces).

There is a clear resemblance between DCSS and PLPA and they are both based on the use of a number of common page tables.

Each user (address space) has its own segment table in which the individual entries point to page tables containing information about the virtual-to-real memory mapping. For VM and MVS/370, each segment entry and its related page table describes 64K of storage. For MVS/XA, each entry describes 1M.

The segment table closely maps the well organized view we normally have of virtual storage (see Figure 1), while the page tables containing pointers to real storage have a much more complex look that even changes over time (see Figure 2).

FIGURE 1

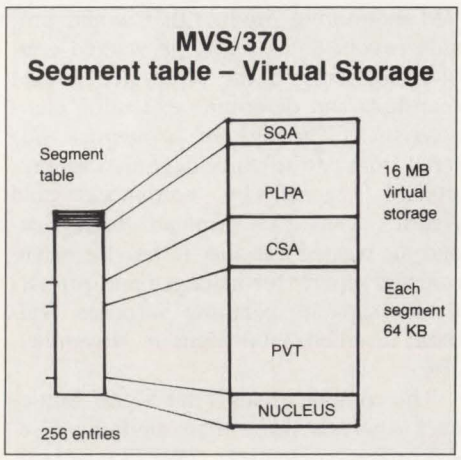

The page tables for the private area (PVT) are only addressed by one single segment table. Page tables for all the common areas (among them PLPA) are addressed by all segment tables in a MVS system.

In MVS, the common segment entries show a static picture in which all common entries are set up at IPL time and copied every time as a new address space; thereby, a new segment table is created.

For VM the situation is quite different. The shared segments are located above the user machine boundary (see Figure 3) and are first addressable following a special LOADSYS function activated through the CP Diagnose interface.

The segment entries above the users machine size are initially set to an "invalid" status and first following a LOADSYS are one or more entries changed to point to common page tables.

The LOADSYS establishes an active bind to a set of programs — all part of a given DCSS. The connection lasts until a PURGESYS function is issued. The LOADSYS and PURGESYS functions are made by normal application programs. The function call specifies a DCSS name as a parameter.

An important difference between VM and MVS is that under VM it is possible to make a bind to DCSSs, overlapping the same virtual storage location if the connections are made at different times. This actually means you can have DCSSA (APL) and DCSS C (CADAM) located at

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

But It *Isn't* Easy!

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

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

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

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

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

BlueLine™
SOFTWARE INC.

New Advances in Performance, Productivity and Planning.

CIRCLE #48 on Reader Service Card ▲

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

FIGURE 2

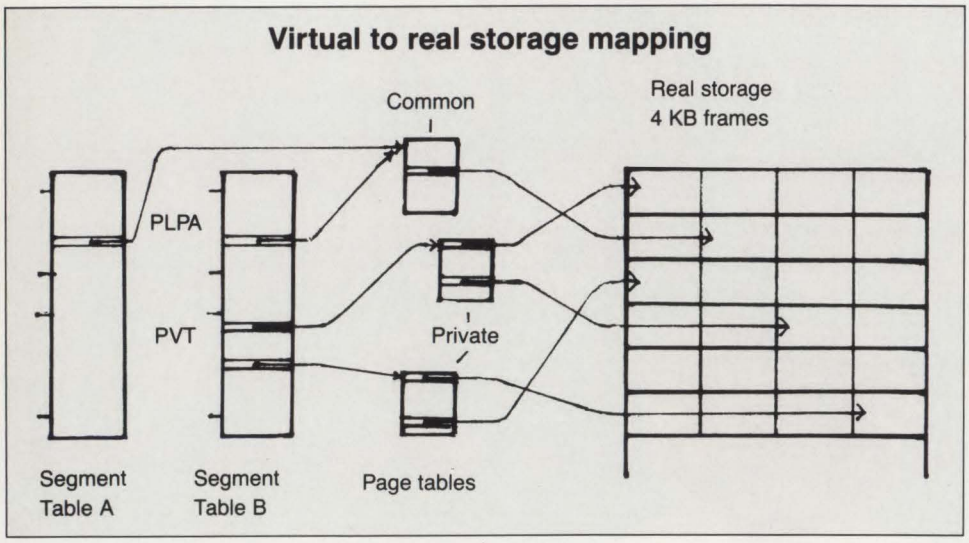

FIGURE 3

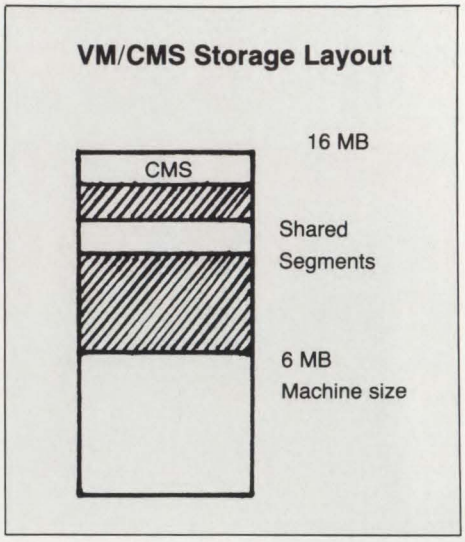

the same storage locations as long as the two applications are not used simultaneously.

Large user machines overlapping one or more DCSSs cannot make reference to these shared segments. However, the virtual storage locations are available for local activities.

In MVS, common code is common for all address spaces and there is a simple one-to-one mapping between common storage locations and the contained program code.

Seen from a maintenance point of view, the PLPA used with MVS is much easier to handle than the DCSSs used with VM. In MVS, all modules are located in SYS1. LPALIB are placed in PLPA at IPL time when the initialization parameter PLPA is specified.

In VM, disk locations must be set aside for each DCSS, the VM csect DMKSNT must be re-assembled when a DCSS entry is added or changed and a new Nucleus must be generated. Further, following an IPL with a changed DCSS layout, a special initial application load procedure must be followed to initialize the page segments with application code.

In MVS, the systems programmer normally has the choice of whether to place a given application in PLPA or not. The decision is normally based on performance criteria. In VM, many CMS applications directly demand the installation in the form of a DCSS.

Normally, programs in PLPA or DCSSs are coded re-enterable; that is, they do not modify themselves. For MVS it is an absolute must. A code update will result in an abend (OC4, segment protection, key zero). For VM, a local copy of the DCSS is created for the user making the update and the segment and page tables are changed accordingly.

The following is a summary about the way VM and MVS have implemented shared segments.

DCSS	PLPA
(+) Dynamic structure, common code can "overlap"	(+) Simple to maintain
(+) Handles non-re-enterable programs	(+) Integrated with normal program load
(−) Difficult to manage	(−) Common means common for all
(−) Application awareness	(−) Static

Most of the problems with DCSSs are related to a lack of high-level control functions to supplement a rather powerful basic design. In the September/October, 1987 issue of *MAINFRAME JOURNAL*, a product from VM Systems Group (V/SEG) was described which tries to overcome most of the DCSS shortcomings.

The basic DCSS design is such that a given strip of virtual storage can contain more application code than just the size of the strip. In MVS terms, the flexibility of DCSSs would be something like a parallel PLPA feature in which an address space dynamically could switch from one to the other.

Thinking back to all the problems with virtual storage for MVS/370, one might ask why some of the ideas from DCSS were not applied to the MVS world. IBM introduced cross memory, but this only resulted in minor PLPA reductions. A DCSS-like feature might have been much

FIGURE 4

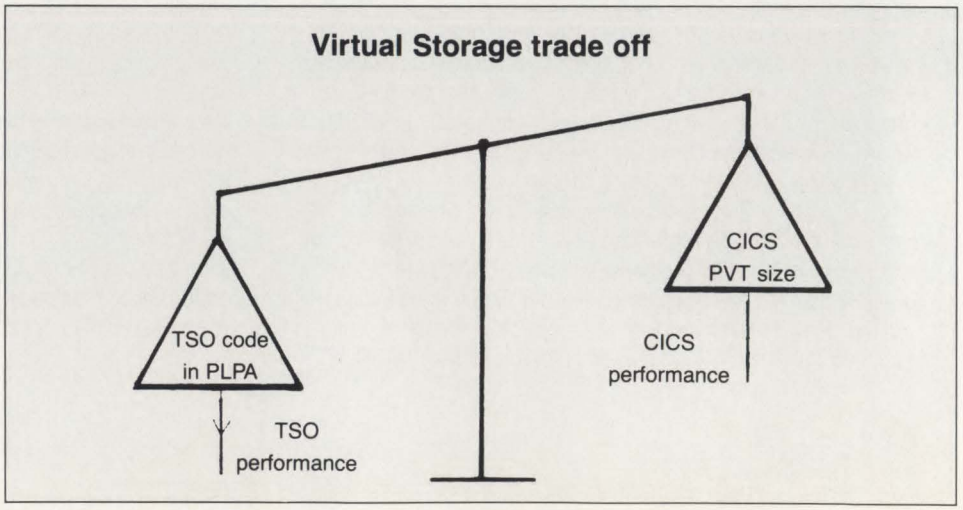

CTOP III™

Reduce CICS Response Time 50-85%

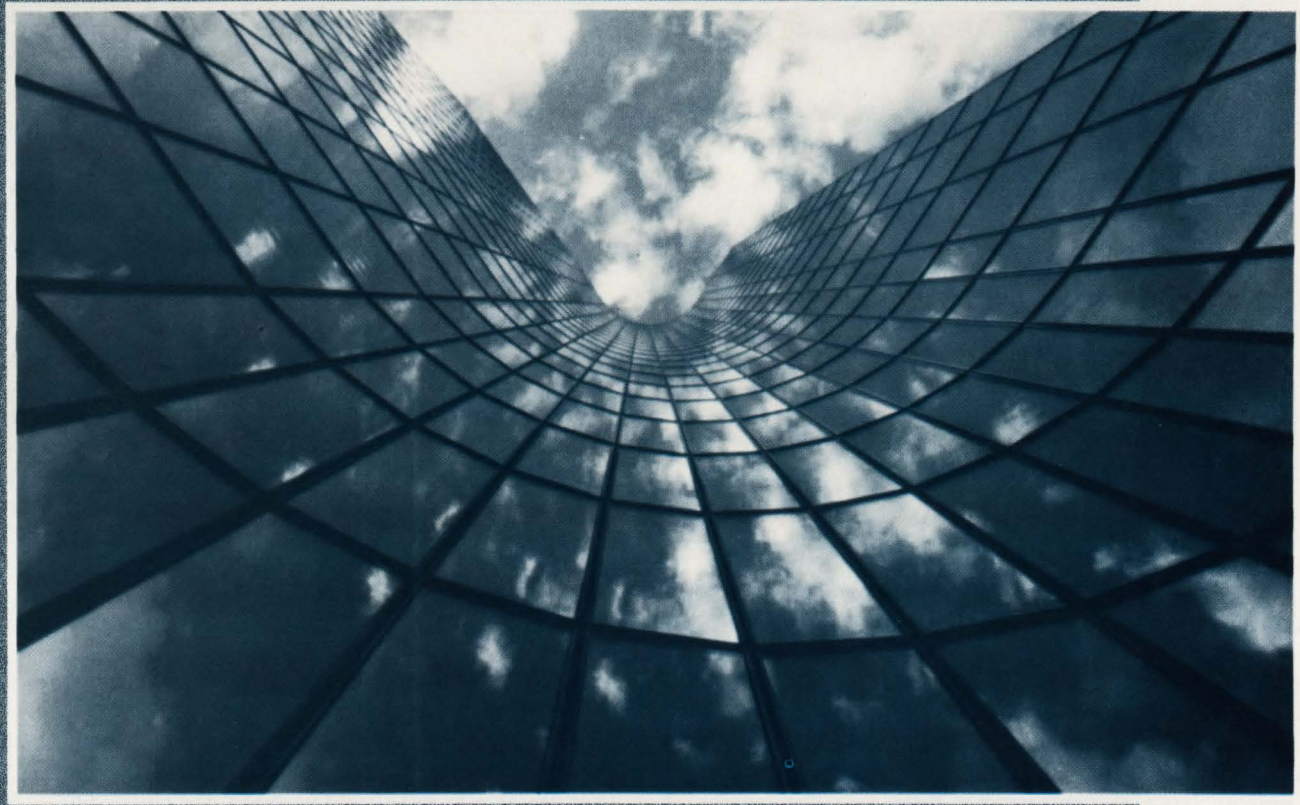

CTOP III™ *"The Proven Alternative"*

- 50% to 85% compression
- $6,650 — DOS/VSE
- $6,650 — OS/MVS — less than 100 terminals
- $16,995 — OS/MVS — greater than 100 terminals

CTOP III™ is a trademark of H&W Computer 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 #89 on Reader Service Card ▲

FIGURE 5

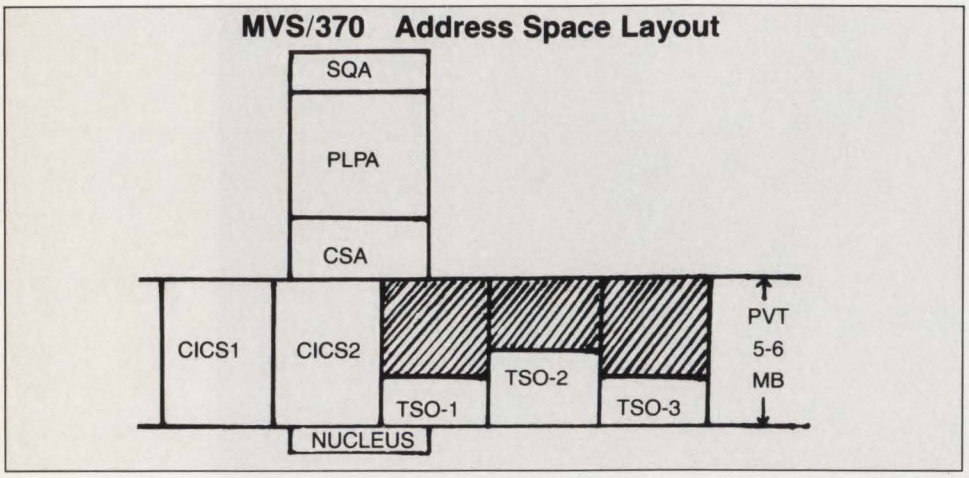

FIGURE 6

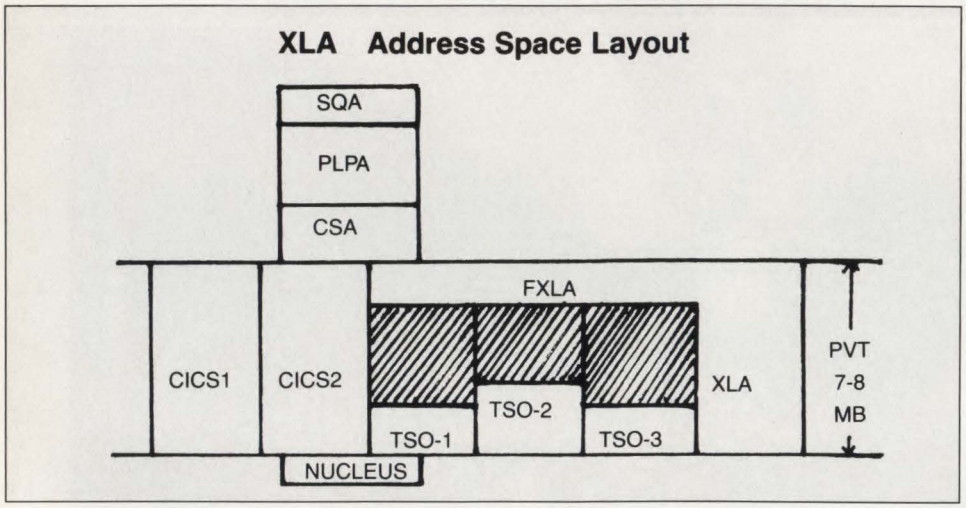

more powerful for a large number of MVS installations.

A couple of years ago, my company, CMA Software, was confronted with a classic virtual storage problem on a MVS/370 system. We were working for an installation in which both CICS and TSO were primary production systems and they wished both to be optimized. The traditional way to tune CICS was to expand buffer pools and the DSA size (to avoid program compression). For TSO, it was to place all the application modules (ISPF, APL, GDDM, commands . . .) in common. The CICS tuning effort would result in an expanded PVT size while the TSO tuning activities would increase the size of PLPA. The problem was that an increase in PLPA automatically resulted in a decrease of the PVT size. We had a situation where tuning one system would damage the other (see Figure 4).

Like many other tuning activities, there was a trade-off between different activities. In this case, the trade-off was created by virtual storage limitations.

The storage layout for the MVS/370 system is shown in Figure 5. Like a lot of similar systems, at least one CICS system was using the private area to the limit while the TSO address spaces all contained lots of unused space. At the same time the situation was that a good part of PLPA (all the TSO related programs) never was referenced by CICS.

The TSO session address spaces have some similarities with VM/CMS machines. With the DCSS concept in mind, in which common storage is not necessarily common for all, we designed storage extension to MVS/370.

The basic idea was to move TSO code from PLPA down to the unused parts of the TSO session address spaces and at the same time to keep the code as shared.

The storage layout for this extended system is shown in Figure 6. The size of the PLPA was reduced and this resulted in an increased PVT size and, thereby, a larger CICS region size.

A new address space, Extended Link Area (XLA), was introduced together with a new type of common area named FXLA. The common area was common for all the TSO session address spaces and the XLA address space.

The functions of the XLA address space were primarily to load all the modules used by the TSO sessions to the FXLA area and to open up the area to be accessible by all starting TSO sessions.

The name FXLA reflects the idea that the common area for simplicity was fixed. We originally designed XLA for a MVS/370 system running under VM. This was just a way to direct the paging control to CP.

The whole setup was later converted to a product called CMA-XLA. CMA-XLA offers virtual storage constraint relief to the MVS/370 world.

In many ways CMA-XLA combines the dynamics and the flexibility of DCSSs with the more manageable PLPA concept. The modules to be part of FXLA are given in a sequential dataset (MLPA list format) which is supplied as input to the XLA address space. The system can be installed and activated without an IPL and no changes need to be applied to MVS system code.

Finally, all modules in FXLA are automatically made known to content supervision, so they are accessed when referenced through LINK, LOAD, XCTL and ATTACH.

For a TSO session to use the FXLA feature, a small change must be applied to the logon procedure. Only sessions using the updated procedure will run with FXLA which of course has special interest when initial testing is performed.

The gains to be achieved from CMA-XLA depend on what has been placed in PLPA in the first place. On the system in which XLA was first installed, we were able to move ISPF/PDF, SDSF and a number of TSO commands from PLPA to FXLA resulting in an increased PVT size of 1.3M. Further, we now could afford to

place APL2 and the GDDM routines used by APL in FXLA so we ended up with a FXLA area of 2.5M.

In general there are a number of other candidates for FXLA (DCF, QMF, CSP, . . .). Everything that can be placed in PLPA is also a candidate for FXLA.

CMA-XLA is not restricted to the CICS/TSO environment. In a shop in which several CICS systems are running in parallel with a storage demanding application and in which this application is the constrained one, CMA-XLA can be used to remove CICS code from PLPA to FXLA, whereby a larger PVT size for the constrained application can be achieved.

In shops with IMS, CMA-XLA can be used to offload PLPA for message region related code (both system and application programs), whereby either the control region or the CSA size can be increased.

The XLA concept could be expanded so more than one XLA system was active in a given system. Again with an analogy to DCSS, the different XLA systems could each control separate common areas, located at overlapping the same virtual storage locations. CMA-XLA only supports one active XLA common area.

MVS/XA has created a more comprehensive solution to the virtual storage limitations, but from time-to-time we still hear about installations with storage problems below the 16M line.

With MVS/XA now getting stabilized, we consider updating XLA so it also works under XA. The page segments used with XA are larger than the ones used with MVS/370. However, the fundamental structures of the paging system (especially for storage below 16M) have not changed.

Once again we might be able to offer an alternative solution to the never-ending storage constraint problem.

Editors Note: CMA-XLA is marketed in the USA by Maersk Data (USA) Inc., Madison, NJ, and by Symark International, Westlake Village, CA.

ABOUT THE AUTHOR

Ole Reitzel Jensen has been working with IBM operating systems for more than 15 years. Holding a Master of Science in Electrical Engineering, Jensen currently works for the CMA Software A/S as an operating system technical consultant responsible for product development. CMA Software A/S, Marielundvej 46B, DK-2730 Herlev, Denmark.

Relational Technology... *Is It For You?*

By John E. Fair

During the last several years, a great deal of emphasis has been placed upon relational database technology. Various vendors have stressed the importance of converting current applications to the new technology. Enterprises have, in some instances, decided that all new database applications will exploit the new products.

Is it wise to arbitrarily decide to use relational database products for all new applications? Do we know enough about the effects of the new products? Have the products reached the same level of reliability as products currently being used? Have the effects of using the products yet been quantified from a performance and capacity planning standpoint?

This article attempts to answer some of these questions by providing historical information about the technology and some of the products and by pointing out some of the pitfalls that arise as a result of poor planning and use. While the article is not exclusively dedicated to DB2, examples used during the presentation are in IBM's SQL/DB2 format.

During the decade of the 70s, the computing power of the average data processing center increased substantially. The complexity of the average database application has also increased at approximately the same rate.

More importantly, the cost of designing, writing and maintaining application systems, the costs of people, have increased at an even faster rate.

During the 1970s and early 1980s, IBM developed two "relational" database systems: SQL/DS for the VM environment and DB2 for the MVS site.

A relational database is perceived by its users as a collection of tables and nothing but tables. Unlike DL/1, in which data is perceived to be stored in a hierarchical manner, relational database files are sequentially organized and usually contain fixed-length records. There is no need for an analyst to be concerned about pointers to records or record segments in another database. Whether or not individual relational database products use pointers is a system design issue and is of no concern to the application developer. Relationships are established by having more than one table with key columns containing the same data values.

A relational database has three major characteristics that are required to satisfy the relational model: a tabular structure; a supporting language (that is, SQL); and data integrity.

Tabular databases are not necessarily relational. In order for a database to be relational, it must have a supporting language and satisfy the data integrity issues contained in the relational model.

Relational databases should have a language that provides mechanisms for defining, manipulating and controlling data and access to data.

Characteristics of a field are specified when the table is defined and data in violation of the basic characteristics is not allowed. Issues relating to data integrity will be discussed in detail later in this article.

A major advantage to a relational database is its simplicity. The programmer no longer needs to know the physical structure of the data; nor does he need to concentrate on indexes for performance. He no longer has to know how to navigate through the data to obtain one or more pieces of data.

With most relational database products, execution of a SQL statement results in many discrete events taking place. One event may be that the system determines whether an indexed search of the data is warranted or if the SET request would best be satisfied by a data scan.

Another event may be that if an indexed search is warranted, the system chooses the proper index to use — if indexes to the data exist. Or, the system may decide if the query can be satisfied by only a search of the index or if the data tables must be read. In many instances, a query can be satisfied by only searching index tables. Further, if a GROUP BY or ORDER BY sub-command is included in the SQL request, the system performs a sort of the data, if required.

It is important that good standards exist for development of any database system. With relational systems, good standards are essential. With these systems, the very parameters that make the systems powerful may also have negative impacts upon performance. During times when the system is very active, it is probably not desirable for an interactive query to take the time to sort a table: an action that would be performed if the SQL statement includes an ORDER BY sub-command.

SET theory processing endemic to all true relational database systems allows users to retrieve or update multiple rows of database tables with a single program

How's obscurantism doing?

Christopher Columbus' caravels were expected to sink into the abysses of the Ocean. The Brooklyn Bridge could not but collapse the very day it would be inaugurated. Then there are those irrational fears throughout History.

Christopher Columbus discovered America, and Brooklyn Bridge spans East River as high as ever.

Who's afraid of a mass migration to MVS?

Mass migration is now discussed alongside traditional piecemeal migrations. It is hitting the headlines. So much for taboos! Comparison has become inevitable.

You'll want to know exactly what to compare:

Lead Times

Proponents of piecemeal migration admit that lead times can range anywhere from six months to two years. The more genuine among them add that another six months are required for preparation and planning.

All mass migrations are achieved within a six to nine months timeframe that begins with preparation and planning, and completes with the "switchover". Mass operations considerably reduce planning and follow-up workloads.

Costs

There again those first six months of preparation seem to have vanished from the minds of most advocators of slow migration. They are just as forgetful of those costs production generates under coexistent VSE and MVS systems, and of the costs of duplicate maintenance.

All things considered, mass migration is between two and three times less expensive. This is tried and tested.

Safety

Trauma sticks to piecemeal migration as spots to dice. Although traditional methods endeavor to free you from the spell, trauma spares no one: the Migration Team, Production, the Company itself; all are hard hit.

With piecemeal migration, the error rate gets appalling. Duplication is a mess for naming conventions, files, programs, maintenance, production.

As if half of the Nation were driving on the left and the other half on the right.

Mass migration neither adversely impacts nor ever delays development and production. You'll be holding your breath the week-end the switchover is triggered, and maybe the next few days. That's right. But that will be only because of a possible hardware failure. All the time you'll know the antidote is available. You can always fall back smoothly upon your regular VSE production.

Some vendors will tell you they tried their hands at mass migration years ago, then gave up. No wonder. They lacked for a codified method and appropriate tools. Only these can keep any bridge from falling down.

Prescriptibility

There are supporters of piecemeal migration shrewd enough to concede advantages to both methods. They say only high integration of systems warrants that a site should migrate "en masse".

Is any DP center today without a data base or TP constraints?

If great tasks can be tackled, so can smaller ones. The reverse does not hold. The piecemeal approach cannot handle highly integrated sites.

Mass migration handles any kind of sites.

CORTEX brings you added value!

CORTEX is the name of a software product and a methodology for mass migration.

CORTEX is the name of a model for automating production under MVS.

CORTEX fully automates translation of:

- Programs from all languages
- JCLs and related materials
- Files

CORTEX makes your VSE shop switch over unscathed to MVS in a week-end.

CORTEX automates your MVS production, thus achieving quality assurance and minimizing staff requirements.

CORTEX is proven.

Wladimir WACHSPRESS
President of SISRO
Co-author of CORTEX

In the United States, only **COMPUTER TASK GROUP INC.** is licensed to offer the **CORTEX** software and services.

CTG 3095 Union Rd.,
Orchard Park,
New York 14127-1214
Phone (716) 674-9310

statement instead of having to process records one at a time.

The relational model, along with its supporting mathematics, was developed by Dr. E.F. (''Ted'') Codd, now associated with the Codd Date Consulting Group. It was first published in 1968 as ''A Relational Model of Data For Large Shared Data Banks.''

Another Twist — Distributed Data

With departmental computing, there is a real need to distribute data to the processing nodes. There are multiple requirements for distributing data.

In the IBM community there have been several attempts to distribute hierarchic databases. IBM currently has no products commercially available that allow for distributing DL/1 data to outlying processing nodes. Other vendors have been slightly more successful, but they have been so on smaller databases than the average DL/1 database.

Relational technology lends itself to distribution of data. Rather than having a data structure maintained by direct pointers, relational data is related by symbolic keys. A field in one table will contain a value (that is, a foreign key) that matches its counterpart in another table. We will discuss these issues in a great deal more detail in the section entitled Relational Database Keys.

There are obvious maintenance, security, reliability and recovery issues associated with distributing any type of database.

Digital Equipment Company (DEC) has a long history for distributing processing, workloads and data. On the other hand, neither IBM nor its customer base have yet accomplished distribution of data on a large scale. This is for a variety of reasons, the most likely being the types and scales of workloads and systems in most large IBM installations. Other vendors and users have also tried to distribute databases, but with mixed success.

Relational technology is the first real hope for distributing databases on a large scale. The elimination of direct pointers from database datasets will reduce the efforts involved in completing this complex task.

There are at least two types of distributed databases, one significantly more complex than the other.

Distribution by Key Range

In the first instance, identically structured files are maintained on more than one processing node. Each processing node contains only the records it needs to perform ''local'' processes. For instance, records for customers of a particular store are kept and maintained by that store's computer. Access to that data from another processing node requires transmitting a transaction requesting data from the requesting node to the processing node with the answer back to the requester.

The type of processing described above exists today in many enterprises. In fact, most commercial DB/DC systems provide some sort of facility for performing that type of function.

Distribution of Data Content

This second type of distributed data is far more complex than the first. It involves separation of data contained in a single database record across multiple processing nodes. As an example, assume that the employee database contained three tables that were geographically separated:

- Employee history data — Corporate — St. Louis
- Employee payroll data — Salary Admin — New York
- Employee skills data — Resource Mgmt — LA

A query against this database could require access to data found in three different geographic areas. While this type of database structure is probably technically feasible, there are significant technical problems involved. These problems include the following:

- Response time problems
- Maintenance and recovery
- Synchronization of data
- Availability of data.

IBM's R-Star (R*) project is currently testing some of the techniques for geographically distributing relational data.

Relational Database

Relational database technology brings with it a symphony of new terminology. Gone are the records and files of yesteryear. The beloved segments, pointers and logical relationships emblazoned on the foreheads of devoted DL/1 users have been put aside forever.

Instead, files, also known as tables, are defined as *relations*. All records or rows in the table have the same characteristics as all other rows in the table. A row or *tuple* is the smallest atomic unit that can be addressed by way of a SQL call. Tuples are made up of fields or columns called *attributes*.

Relations contain several properties unique to relational database systems. One is that within a single base relation, there are no duplicate rows. All primary keys must be unique. Likewise, there are no duplicate column names. The order of both rows and columns is insignificant.

Another is that between tables, no structural links such as direct pointers are visible to the user. Instead, association between the tables is achieved by use of comparison of values in specified columns of the specified relations.

Domain is another term used extensively in relational database products. A domain limits the range of values that an attribute can contain. Better stated, a domain is a pool of values from which an attribute's value is extracted. For instance, an ''infinite'' domain defined as character format may contain any character values, without restriction, whereas the ''finite'' domain SEX may include only the values ''M'' and ''F.'' The domain for the attribute STATE will contain 50 values, a value for each state in the United States.

Keys

Even though a relation is defined as having specific attributes and an unspecified and unordered number of rows, data in the relation is usually logically ordered by some attribute known as the *primary key*. Even though data is stored in unordered format, an index usually exists to allow ordered retrieval by the value of the primary key. A basic premise of relational technology can be stated as follows:

> ''Information in a relation is unchanged if rows are reordered. Likewise, information in the relation is unchanged if attributes, including column headings, are reordered.'' (E.F. Codd, ''The Relational Model for Formatted Databases,'' 1985)

Order of data is achieved by having a *cluster index* on the primary key to ensure that queries will not result in frequent, random access to the tables. A relation may have only one cluster index.

Consider the following example. A table exists with the following organization:

```
(EMPNO     : 1234)
(LASTNAME  : 'Smith')
(CITY      : 'Vienna')
(STATE     : 'VA')
(ZIP       : '22180')
```

The primary key for this table is used to ensure uniqueness of each tuple within the relation. In order to ensure that all primary keys are unique, no component

of a primary key is allowed to contain NULL values. For this example we assumed that the table has two *candidate keys:* EMPNO and LASTNAME. For this example we will choose EMPNO as the primary key with LASTNAME as an alternate key because of the probability of lack of uniqueness should LASTNAME be chosen as a primary key. LASTNAME then becomes an alternate key.

Another type of key common to relational database systems is the *foreign key.* If a relational database contains more than one relation, there must be ways to connect the tables. The foreign key attribute becomes the connector between the relations.

Assume that a second table exists as part of our database. Given that the structure of this table is

(SKILL : "expert")
(E_NO . . . . . : 1234)

the attribute E_NO becomes the "foreign key" to the first table.

The Language

Since one of the major components of a relational database is its language, we would be amiss to omit discussions of the language that is becoming the "de facto standard" for relational processing: the Structured Query Language (SQL).

Even though a standard version of SQL exists, many variants of the language currently exist. These changes are usually enhancements to better support their respective products. IBM's two implementations of the language are not yet identical, even though we understand that work is underway to reconcile the remaining differences.

The SQL language contains three types of statements:

- Data definition statements
- Data manipulation statements
- Data control statements.

With SQL the user can define relations and table spaces, can access the relations and can control access to the data and to the system.

The language allows the user to select the data to manipulate. It does not give the user the opportunity, nor does the requirement exist, to specify how the data is to be retrieved. Relational systems generally select what they consider to be the best path to data. They also do not usually inform the user of the path selected.

The *SQL EXPLAIN* function is one way to monitor the efficiency of relational databases generated by DB2 and SQL/DS. Among other things, the EXPLAIN function identifies whether an indexed search or a scan of rows of the relation was used. If applicable, the name of the index is also included. EXPLAIN also reveals whether access to the data was necessary or whether the query could be satisfied by reading only the index relations.

The EXPLAIN function is so important that many installations have developed tight standards requiring its use.

Assume that we have a table named *Employee* that contains three attributes — NAME, SKILL, EMPNO. The table looks like this.

Name	Skill	Empno
Adams, T.	Programmer	1234
Moore, C.	Operator	0840
Smith, D.	Manager	1250
Young, M.	Analyst	1300

Also assume that the database contains a second table named *Skills* that contains the following attributes: SKILLS, DESC, BASESAL.

The inquiry SELECT * FROM EMPLOYEE WHERE
SKILL = 'Operator';
will return all data about C. Moore.

The inquiry SELECT * FROM EMPLOYEE, SKILLS WHERE
SKILL = 'Programmer';
will return T. Adams' data from the EMPLOYEE table and attributes DESC and BASESAL from the SKILLS table. The attribute SKILL is the foreign key used to query the table SKILLS.

Referential Integrity

"If base relation R2 includes a foreign key FK matching the primary key PK of some base relation R1, then every value of FK in R2 must either be equal to the value of PK in some tuple of R1 or be wholly null (that is, each attribute value participating in that FK value must be null)."

All foreign keys must have attributes with matching primary key values in some tuple of the relation. The only exception to the rule is when the foreign key contains null values.

One weakness of many existing relational database products today is that they generally provide no mechanism for ensuring that the "referential integrity" rules included with Dr. Codd's relational model

are enforced. IBM has committed to providing full referential integrity with some future release of their DB2 product.

PROS of Relational Database

Relational database products have greatly simplified the development cycle for many new applications. Even though programmers must scale a learning curve, the learning cycle is usually short enough that it has little, if any, impact on application development. How much impact the learning curve will have depends upon a number of factors:

- The project schedule (how much time is there?)
- The size of the project
- The complexity of the application
- The criticality of the system.

Most vendors, including IBM, recommend that the first application using their relational products be relatively small, relatively simple and on a non-critical development cycle. This reduces many of the pressures typically attendant to project development.

It is also wise to construct a test or prototype system to measure activity and effects of the new system. Unfortunately, prototype systems typically give little information about how the system will respond under load and when databases become very large. A system of this type will, however, identify basic design flaws that may introduce performance problems into the system. The prototype system will also identify SQL coding conventions that will have far ranging negative impacts on transaction response time.

It is essential that applications running on prototype systems use the diagnostic tools available including the EXPLAIN function that will be described later in this article. It is also wise to have some type of monitor that interactively reports workload activities.

Several major issues constitute the major advantages from a developer's point of view. The fact that the developer no longer needs to know the physical structure of the data when combined with the ability to do SET processing eliminates two major bottlenecks to development of hierarchic database systems. With hierarchic systems or flat files, it was essential that the programmer know how the data was physically stored on the media in order to optimize performance and to ensure the proper retrieval of records using "Get Unique" and "Get Next" type processing.

With SET processing, multiple records can be retrieved and/or updated with a single database call, unlike with DL/1, that requires retrieving and updating individual segments.

Data navigation is no longer an issue with relational systems! Since the programmer no longer has to worry about how the data is stored, he has no major concerns over how the data is ordered or structured on the media.

Relational databases should be normalized to at least second normal form. First and second normal forms are described in general terms below. Database designers should be very familiar with normalization of databases. A great deal of information is available to aid database designers.

If a relation has a simple primary key meaning that the key contains a single attribute, any repeating groups of data should be removed from the relation. Since an attribute can contain a single atomic value, repeating groups cannot be allowed. This is known as first normal form.

If the primary key of a relation is a compound key meaning that it includes more than one attribute, columns should be removed from the relation unless they are dependent upon the entire key. Better said, if an attribute does not depend upon the entire key, but rather a subset of the key, it should be removed from the relation and placed into another table. This is known as second normal form.

Most relational systems determine the best path to use for retrieving data from DASD. This eliminates current concerns over "which secondary index do I use." Even though the DBA must still select attributes upon which indexes are to be constructed, the use of the index is of little consideration to the developer.

The net effect of using relational technology is a marked reduction in development time resulting in reduced development costs. Possibly more important than reductions in development time and costs is that the resultant compact and simple systems will also be easier to maintain.

One caveat must accompany the "PROS" for using relational database. In order to gain benefits from relational technology, features of the technology must be exploited. The new application will achieve no gain if coding techniques for hierarchic database systems are replicated into new systems. The real power of the new technology is in its ability to do SET instead of RECORD processing.

CONS of Relational Database

It seems that the hype associated with relational databases would cause most enterprises to flock to jump on the relational bandwagon, and many are. But there are some reasons to wait.

Relational database technology is still relatively new. It has been available for a short time and though the user base is growing very rapidly, it still falls far short of the user base for conventional database systems.

Endemic to most new systems are performance and reliability problems. Most of the commercially available relational database systems, and especially DB2, have exploited the recoverability options of the MVS operating system and have been very reliable. We are led to believe that few users have complained about system reliability. Therefore, we do not consider this to be a serious issue.

Performance has been a serious issue with some of the products. In early versions of the relational products, enterprises found it difficult, if not impossible, to process required transaction volumes using the new technology. Later versions of the products have largely alleviated these problems.

Some rules of thumb still exist for companies considering installing applications using relational database systems.

Most relational database systems are resource intensive. They obtain their processing speeds by keeping as much as possible in memory at all times.

Even though memory is more an issue than cycles, it takes additional cycles to run the address spaces used by the relational products.

Use of relational products is a tradeoff — resources versus people. Organizations placing heavy emphasis upon people costs will probably find that the additional resources are more than justified. Companies trying to reduce hardware costs will almost certainly experience increases in their hardware budgets unless they currently have excess capacity.

Organizations should realize that growth in the numbers of transactions and ad hoc queries will demand that performance of the system be carefully monitored. As system loads increase, it becomes important that companies have carefully planned performance strategies and that companies understand issues that affect performance. An organization should also understand that tuning of relational database systems may require changing the

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!

CIRCLE #32 on Reader Service Card ▲

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

characteristics of the DBMS in addition to those of the applications.

Since relational systems tend by nature to be resource intensive, improvement of bad or marginal performance may involve one or more of the following: addition of real memory; addition of computing power (MIPS); reduction of concurrent workloads; and application tuning.

Some programmers feel ill at ease not looking at every record processed. This is probably a temporary condition that will pass, but it is sometimes a major hurdle to scale.

Some developers are concerned that some of their flexibility has been removed due to the system making more of the decisions. Again, this will probably pass.

Relational databases MUST BE normalized. Redundancies must be eliminated as much as possible. While this was desirable using hierarchic systems, it is ESSENTIAL using relational systems. Again, there is a tradeoff. Users of relational database systems must have DBAs with slightly more expertise than organizations that use flat or hierarchic databases. The tradeoff is: more experienced DBA or more experienced programming staff.

IBM has committed to resolving referential integrity issues with some future release of their DB2 product. Other vendors have similar issues that must also be resolved.

Until the issues are solved, applications utilizing foreign keys should make some provisions to ensure that matching foreign key attributes exist in all participating relations.

Replacement for Current Systems?

The question is frequently asked "Are my current applications candidates for relational technology?" The answer is a resounding "maybe!" Relational database systems are good candidates for many new applications under consideration. It may or may not be a good idea to consider converting existing systems to fit under the new technology. We will discuss some of the types of systems that lend themselves to relational technology products. We will also discuss some of the reasons why the technology should not be seriously considered YET for other types of applications.

Information Centers

The ad hoc inquiry is still the best use for relational database technology. In fact, in the earliest days of the products, vendors openly endorsed using their products for only those types of applications.

For instance, using DB2 with QMF allows managers and "non-data processing people" to build sophisticated queries without knowing programming techniques or languages. Furthermore, SET theory processing allows retrieval, basic statistics on and updates to multiple rows of the relation.

Once it was proven that the products could function well in the information center environment and once performance of the products was determined to be adequate for other environments, the products came into use in other areas.

The promise of portability of relational databases to personal computers and to intelligent workstations also makes relational database products very attractive.

Commercial Applications

Whether or not relational technology is useful to a commercial application will depend upon factors such as: transaction volume; transaction arrival rate; synergy of transactions; transaction profiles (that is, queries, updates, and so on).

Relational technology is clearly not intended for applications requiring transaction volumes in excess of 75 per second. It is also probably not for applications that require that large numbers of users retrieve and update data concurrently.

High Volume Applications

Current releases of relational database products have not yet achieved performance levels to allow applications that currently require IMS Fast Path or equivalent systems.

Steps to Sound Relational Database

Before a database application system can successfully be installed, several steps must be completed. In the following section, some of the issues involved will be identified. Even though these issues relate to relational database systems, they are equally applicable to other types of database application systems.

The database must be designed by a DBA with a thorough understanding of the requirements of the application. Without clear definition of what the application is supposed to do, the designer will have little chance of success.

After acquiring a thorough knowledge of the application, the DBA should develop a sound logical model. The logical model involves identifying all data elements and data sources that will be needed by the application. This collecting of data is without consideration for physical data structures.

Following development of the logical model, two key parts of the project must be completed. These phases may run concurrently, be interactive or sequential; but, there will be a great deal of interaction between these phases. The phases are basic application design and development of the physical database model. Final application design is dependent, to some degree, on the physical database.

Recommendations

The decision to change from a tried and true technology such as DL/1 to a newer and seemingly more radical one like DB2 or Oracle may seem frightening at first. With respect to DB2, it has frequently been said, "It is not a matter of whether we are going to use DB2, but when." This is probably true of relational technology in general.

With this in mind and with the advancements that the products have made since the beginning of 1986, there is probably no time better than NOW to consider using these tools.

Keep in mind, however, that the tools are significant departures from tools of the past. They trade computer resources for people!

Proceed cautiously, but proceed. The tools provide you with great capabilities. There is no better time than the present to start tapping these resources and capabilities.

Do not jump in all at once. Start by developing a small, non-critical application first. Ride the learning curve on an application that will not be in the enterprise's critical path.

Relational systems tend to be resource intensive. Be sure that MIPS and real memory are available to run the systems. If you are considering DB2, do not do it unless you are running MVS/XA.

Do not leave things to chance. Prototype! Test! Model! Monitor!

ABOUT THE AUTHOR

John E. Fair is a senior consultant with Morino Associates. He is the author of more than 20 published articles on IMS and DB/DC systems. His 18 years of experience have been largely in banking, manufacturing and insurance industries. The company is headquartered in Vienna, VA (703) 734-9494.

SYSTEM DEVELOPMENT UNDER CICS

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®

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 ▲

Electronic Coupon Distribution System Supported by IBM 9370 and PCs

By Edith Myers

Coupons offering consumers discounts from the purchase price have been a successful merchandising technique for decades. They have been offered by manufacturers in hopes consumers would choose their products over those of their competitors for savings of anywhere from two to 50 percent. Coupon distribution through the years has been more of a shotgun rather than a rifle approach. Manufacturers could only hope that consumers who received their coupons would be interested in their product in the first place. In the second place, they had to hope the consumers would be sufficiently interested to either clip a coupon from a publication or pick theirs from a cumbersome collection in a direct mail offering.

Now, a sophisticated computer-based system has refined coupon promotion in a way that the first users of the practice would never have dreamed possible. Merchandise manufacturers can target their competitors' products to trigger their own coupons, thereby putting a purchase incentive in the hands of consumers who already have demonstrated an interest in the manufacturer's type of product. Or, they can see that their coupons reach users of products to which theirs are naturally complementary.

All this is made possible by a system called the Coupon Solution that distributes coupons tied to electronic Point-Of-Sale (POS) devices in supermarkets. The company that has made it possible is a fast-growing, venture capital-backed, Anaheim, CA startup — Catalina Marketing Corporation.

An Idea Comes to Life

Founded in late 1983, Catalina's founders came out of package goods marketing, supermarket retailing and data processing. The original idea for this unique concept came from Tom Mindrum, senior vice president of marketing. During 1984 and 1985, he and the other four founders refined the idea and established a thriving business based upon the Coupon Solution system. Its methods of electronic coupon distribution at the point of sale are protected by patents, not tied to either the hardware or software involved.

Like most innovative ideas, the concept for the Coupon Solution is rather simple: instantly printing coupons based upon the customer's actual purchases. The implementation of the concept, however, is quite complex due to the variety of supermarket POS equipment and the difficulty of integrating readily available components to produce the desired results. Catalina, with the help of outside consultants, solved these problems by using a combination of off-the-shelf products and custom interfaces. The idea was to deploy in the stores PCs that utilize a proprietary hardware interface to the stores' installed POS systems and print coupons on specially designed printers located at the checkout stand. The store level system was supported by a central host system and host-to-store communications. Development of the system continued through 1985.

The aim of the development efforts was to build a system that would be compatible with the majority of installed scanning POS systems. The system today is compatible with the leaders: IBM's 3660, 3650 and 4680 systems and National Semiconductor Corporation's Datachecker 1600 and 1700. Nearing completion is an interface to NCR's POS systems.

Today, the Coupon Solution system features a three-tiered architecture: an IBM 9370 Model 60 for order entry; a PC/AT-based communications system in regional offices for transmission of information to and from the stores; and an in-store system including IBM PC/ATs (or clones), plus a special hardware board and specially designed printers at checkout stands.

Customer Rollout

"The first hurdle was selling the concept to retailers. We couldn't interest manufacturers until we had a reasonable number of stores signed up," Catalina's President Michael O'Brien points out.

Catalina signed up its first major chain, Ralph's Grocery Company in Southern California, in 1986. "After that, manufacturers were quick to test," O'Brien states. By the end of 1986, there were 200 participating stores. This grew to 400 by the end of 1987 and today there are 1,500 with more expected to sign up before the end of this year.

There are several incentives for the stores. They are paid per coupon issued

by Catalina. In addition to this, they receive eight cents per redeemed coupon by the manufacturers over the coupon's face value. Then, there is the ability to use the coupon-generating system for their own purposes.

O'Brien notes that some 220 billion coupons are issued (from all sources) per year or 2,800 per household. Of these, approximately seven billion are redeemed. He said Catalina's coupons have a far greater redemption rate because they hit ready targets and generally offer a greater discount, usually 50 cents or higher. Also, the redemption rate is closely tied to the amount of discount.

Evolution of the Host System

As its business has grown, so has Catalina's host computer system. Initially implemented on a Wang VS, it ran on this host for four years. In 1987, Catalina decided an upgrade was necessary. The IBM 9370 was selected as the host computer and the application software was redesigned to take advantage of new application features the 9370 made available. The new system went live on the 9370 in April of this year.

Unicorn Systems Company, a Los Angeles-based consulting and software products company, has been involved with the Catalina operation since the early development stages. Involvement included both consulting and providing software development and execution tools. Unicorn's products allow Catalina's 9370 to function like a large mainframe without requiring the typical high cost, resource intensive data processing environment.

From the beginning, the host system was built to be CICS-based using Unicorn's proprietary CICS emulation products. These products permit CICS to run on desktop microcomputers or under VM on any system utilizing IBM 370 architecture. The VMCICS product family, released in mid-1987, was the first to take CICS into the VM world. This family of products includes VMCICS/Development System (VMCICS/DS), VMCICS/Execution System (VMCICS/ES) and VMVSAM. Together, these three products provide a comprehensive development and execution environment for CICS/VSAM applications directly under the VM operating system.

The availability of these products contributed to the choice of a 9370/60 as host processor for Coupon Solution but there were other reasons as well. Tim Cherney, vice president of MIS for Catalina, says, "I wanted something with 370 architecture because that would make it easy to find both people to work with the system and software packages to run on it."

VM was selected as the host operating system, partly because it seems evident that this is a strategic operating system for IBM. Additionally, it is highly interactive, cost effective and easy to support.

VMCICS and VMVSAM were the logical choices for system software. A primary reason was their ability to run native CICS under VM with good response time. In addition, use of these products minimized Catalina's licensing fees for systems software.

Managing Coupon Orders

The nerve center of Coupon Solution is the on-line Order Control System (OCS) that runs on the 9370 under VMCICS/ES. This system was developed using

VMCICS/DS which allowed two programmers to develop and test 125 CICS programs in about four months. Catalina used its VMCICS/DS screen-painting facility to develop new screens. The company is continuing to use it to maintain screens and create additional screens as applications either change or are added — both regular occurrences with Coupon Solution.

The OCS is the central point for entering and processing coupon orders from manufacturers. As orders are received from the field sales offices, client service personnel enter them into the OCS. Orders are typically for a particular coupon offer that is valid for a period of time (a cycle) in one or more geographic market areas. Powerful facilities allow client service to copy orders from one store to another, one chain to another or to all stores, chains and markets. Once orders are entered and verified, they are ready for transfer to the stores via the communications system. The system has both foreground and background transactions with equivalent daily transaction volumes in excess of 15,000.

All CICS commands in the OCS application are embedded in two special routines called the Screen Interface Routine (SIR) and the File Interface Routine (FIR). SIR and FIR are tools that insulate applications programmers from the CICS command language and afford a degree of system portability by containing all CICS-specific code in centralized routines. "The SIR and FIR software was especially valuable to Catalina since it allowed me to recruit and train COBOL application specialists rather than CICS experts," declares Cherney.

Cherney's data processing staff today numbers 10 people: three in operations, three mainframe applications programmers, one system administrator and three microcomputer programmers who develop applications for stores.

Positioning for Future Growth

From less than a half dozen employees at its inception, Catalina has grown to 95 employees today. It is serving 16 market chains in 14 geographic areas. This represents a potential exposure for clients to 25 million shoppers per week. Catalina's manufacturing clients today number 156. Started in quarters shared with another company in Los Angeles, the firm expanded to new headquarters in Anaheim, CA in May of this year. It also maintains nine regional offices throughout the U.S.

Catalina has been adding stores at the rate of about 100 per month in 1988 which means the system demands are constantly growing. In order to contend with this growth, Cherney anticipates upgrading the host hardware to a 9370/90 next year, "I've put it in my '89 budget." He adds that the system as it exists today has capabilities not yet being utilized.

"The IBM 9370 and VMCICS were the most cost-effective data processing solutions available to us when we planned our system. In retrospect, they were wise business solutions as well," concludes O'Brien.

ABOUT THE AUTHOR

Edith Myers is a free-lance writer based in Los Angeles, CA. She was the Los Angeles bureau manager for Datamation for more than 17 years.

MVS Cross Memory Services

Protection and Authorization

By Bill Carico

Author's note: This article assumes the reader is familiar with the concepts and terminology presented in the first article in this series on Cross Memory Services that appeared in the last issue of MAINFRAME JOURNAL.

Trying to write about how Cross Memory Services (XMS) works is a significant undertaking that could go on almost indefinitely. To keep the task manageable, the remaining two articles in this series will target two technical areas that are relatively complex and easily misunderstood. This article will explain how XMS protection and authorization mechanisms work. The third and final article in this series will look at the macros used to establish the Program Call environment and examine in detail the linkage tables used during the execution of the PC instruction.

Getting Started

Cross Memory Services (XMS) provides a group of macros that may be used to create a cross memory environment and establish the necessary authorization levels and linkage tables. The provider of a cross memory service, usually an MVS subsystem, is responsible for initializing these structures by using the macro interface. Once the appropriate structures are in place, eligible users issue cross memory instructions to access services provided by the subsystem. For this discussion, the term "subsystem" will be used in a broad sense to describe any system component or vendor product that is providing services to a number of users running under MVS. As mentioned in the previous article, this includes PC/AUTH (the XMS supervisory address space), Global Resource Serialization (the ENQ, DEQ, RESERVE functions), TRACE (which is the system trace table), Allocation, CONSOLE (which handles operator communications), LLA (Linklist Look-Aside — contains PDS directories for LNKLST libraries), System Management Facility (SMF), CATALOG and DUMP Services. Examples of optional subsystems using XMS would include products such as IMS, DB2, and many other vendor products such as Candle Corp.'s (Los Angeles, CA) Omegamon.

FIGURE 1

Semi-privileged Instructions used with XMS

EPAR	–	Extract Primary ASID
ESAR	–	Extract Secondary ASID
IAC	–	Insert Address Space Control
IVSK	–	Insert Virtual Storage Key
MVCK	–	Move Characters with Key
MVCP	–	Move to Primary
MVCS	–	Move to Secondary
PC	–	Program Call
PT	–	Program Transfer
SAC	–	Set Address Space Control
SPKA	–	Set PSW Key from Address
SSAR	–	Set Secondary ASID (load CR 7)

The PC/AUTH address space contains the programs that are executed and the control block structures that are built as a result of a subsystem defining the cross memory environment. These services are invoked when the subsystem issues the XMS macros to specify authorization levels and to build linkage tables. The authorization levels that are established serve to control which users can access the subsystem's services. The linkage tables that are built support the special cross memory instructions that were added to the System/370 instruction set.

Semi-Privileged Instructions

As an important part of early system design, certain instructions in the instruction set that modify important system controls are reserved for use by the operating system and its components. These restricted use instructions are called *privileged instructions* and may be executed

FIGURE 2

PSW Key Mask

- Extension to System/370 Key
- Provides tasks and SRBs with a list of keys
- Used for problem state, not used in supervisor state
- 16-bit string loaded into CR 3

Example *This shows the PKM for a key 8 user:*

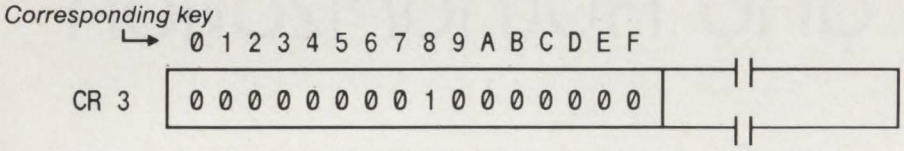

0 - *unauthorized to use key*
1 - *authorized to use key*

only when a program is running in supervisor state. Supervisor state is established whenever MVS sets bit 15 of the current Program Status Word (PSW) to zero. This occurs most often as a result of an interrupt whereby all new PSWs loaded already have bit 15 set to zero. When the system dispatcher loads a PSW to give control to a user program, bit 15 of the PSW is set to one which keeps the user programs in submission. This is known as problem program state, or simply *problem state*. If a problem program issues a privileged instruction, a program check interrupt (privileged operation exception) occurs. Thus, MVS can maintain system integrity and keep users from taking over control of the machine.

Quick VM Tangent

Just as a point of interest for VM users: VM runs its guest machines such as DOS and MVS in problem state. This means that the guest operating system causes a program check interrupt each time it tries to enter supervisor state and execute privileged instructions. Therefore, when the resultant program check interrupts occur, VM regains control and ensures that the desired functions are performed according to and consistent with VM's requirements. This allows VM to maintain control over the environment and also explains some of the overhead which is attributable to VM design.

Back to XMS

A new series of machine instructions was introduced to support XMS. These instructions may be executed in problem state, but will only work correctly if certain authorization checks are successfully completed. Therefore, the category "semi-privileged" was added to describe this new classification of instructions. Authorization checking varies by instruction and is performed by the hardware.

The PC instruction is an example of a semi-privileged instruction. Other XMS instructions that are semi-privileged include the Program Transfer (PT) and the Set Secondary Address Register (SSAR). For a complete list of semi-privileged instructions used with XMS refer to Figure 1.

PSW Key Mask

When XMS was introduced, an extension was made to the System/370 key such that a unit of work (TCB or SRB) is given a list of keys that it is authorized to use.

No special authorization is needed for a program to switch to a key that is in its list. The list of keys is represented by a 16-bit string value called the PSW Key Mask (PKM). Initially all programs are dispatched with a PKM value equal to the bit mask representation in the TCB or SRB (in the TCBPKF or SRBPKF fields respectively).

A user's PKM is loaded into Control Register 3 each time a program executes. See Figure 2. The 16 bits in the PKM correspond positionally to the 16 physical protect keys available in the hardware. The PKM is used to check the authority of programs in problem program state but is not used for programs in supervisor state. The value may be changed as the result of a PC or PT instruction and by the MODESET SVC. In order to use MODESET, a program must be authorized under the Authorized Program Facility (APF). If an unauthorized program issues MODESET, an abend will occur.

Do not confuse XMS authorities with the APF provided by MVS. APF is a separate facility which basically allows selected users to switch from problem state to supervisor state.

The Set PSW Key from Address (SPKA) instruction existed before XMS but as a privileged instruction. With XMS it was changed to semi-privileged status to allow its use from within programs operating in problem state. The instruction works the same as before when executed from supervisor state. From problem state, the execution of SPKA is now controlled by the PKM. This allows a problem state program to set the PSW Key as long as the bit in the PKM corresponding to the Key value is one. Otherwise, the instruction is suppressed and a program check interrupt (privileged operation exception) occurs.

The PKM is also updated, potentially, by the PC and PT instructions. Since a called program may need to operate from a key that is different from that of the calling program, the PC instruction includes logic to update the PKM. Each program invoked using a PC instruction has been previously defined in a table called the Entry Table (ET). Each Entry Table Entry (ETE) contains, among other things, a field called the Execution Key Mask (EKM). As part of the PC instruction logic in microcode, the EKM is OR'ed (combined) with the caller's PKM to produce a new PKM for the called program. When the called program completes, it issues the PT instruction to return control to the caller which in turn changes the PKM back to its original value.

While in cross memory mode, a program may use the MVCP, MVCS and/or the MVCK instructions to move data. An additional access key is supported and is specified as part of these instructions since the source and target locations for the move may be in different protect keys. This second key is in addition to and will probably be different from the program's key contained in the PSW.

The hardware uses the PSW key to control access to one operand (whether it is the source location or target location differs based on the type of instruction) and uses the key specified in the instruction to control access to the other loca-

tion. Naturally, when using MVSCP, MVCS or MVCK in problem state, the PKM is used to control the key values that may be specified from within the instruction as the additional access key.

Authorization Indexes

Unlike supervisor state which provides blanket authorization to its users, XMS has more sophisticated control mechanisms in place to restrict user access. Permission to gain cross memory access may optionally be restricted to selected address spaces. Stated another way, different subsystems can support their own community of address spaces and *not* be exposed to unauthorized access from other users of XMS. In addition, the design of the linkage tables for the PC instruction provide the ability for a subsystem to maintain control over access to selected sets of programs. On the other hand, a subsystem may be a provider of global system services that should be made available system-wide to all users. XMS is set up to handle either requirement according to the needs of the subsystem.

XMS provides two levels of address space access authority: primary and secondary. Primary authority allows a program to issue the program transfer (PT) instruction to return to an address space. This is a confusing area to explain without knowing much more detail than has already been presented. As a simplified explanation, be aware that part of the security for a PC/PT instruction sequence was implemented on the return leg of the journey — the PT portion. The PC authorization checking could theoretically be circumvented, although it is unlikely that this would occur. However, to make sure the security is iron clad, a specific check is made during the return step (the PT instruction) which serves as a double-check of the user's authority to have issued the PC in the first place.

Secondary authority permits one address space to gain secondary addressability to another using the SSAR instruction. Having secondary authority allows a program to move data between the two address spaces.

Each address space using XMS is assigned an Authorization Index (AX) value. A program runs with the AX of the primary address space. Furthermore, each address space is connected to an Authorization Table (AT) which contains the AX definitions. The AT consists of consecutive two-bit long entries that show the two levels of authority for each AX that is in use.

An address space's authority to access another address space is based on the bit settings for its AX in the corresponding AT entry. However, the target address space's AT must also have the correct bit settings for the same AX before a cross memory environment may be established. A program's AX is used to index into the target address space's AT on a PT and SSAR instruction.

In Figure 3, a non-zero value for "P" allows a program to access the AT owner's address space using the PT instruction. A non-zero value for "S" allows a program to gain secondary addressability to the AT owner's address space using the SSAR instruction. In the example in Figure 3, AX-1 and AX-2 have both primary and secondary authority (both bits of the entry are one) while all other AXs have no authority (both bits of the entry are

"Ready to get rid of your image copy problems?"

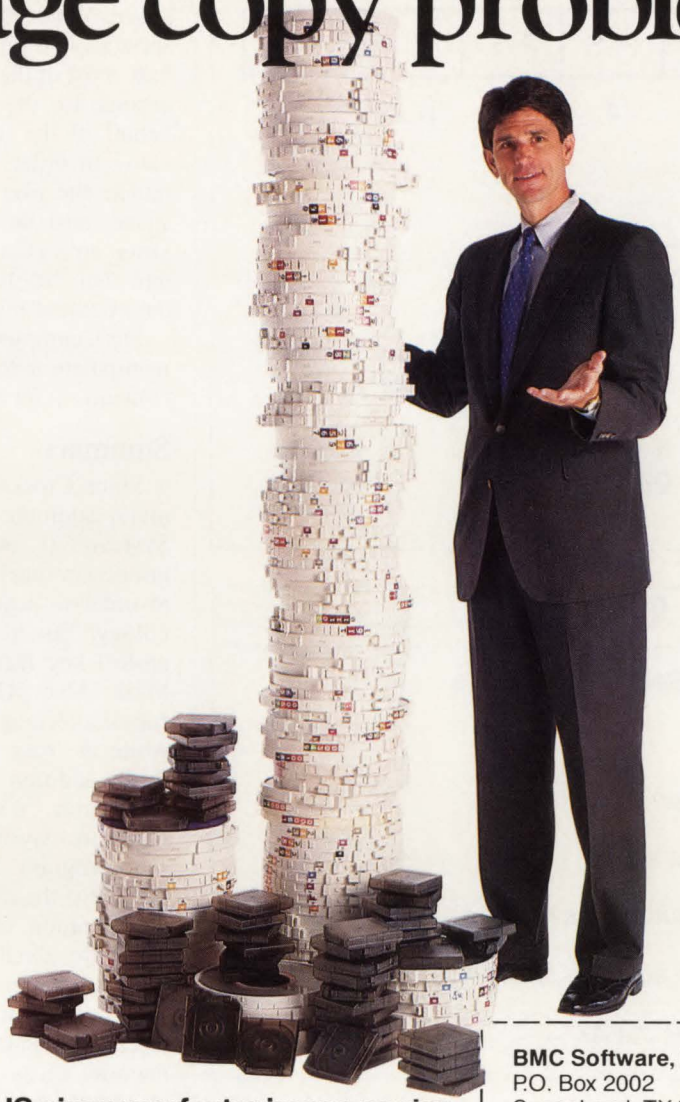

IMAGE COPY PLUS gives you faster image copies with little or no data base down time.

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 create 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.

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 and Canada or 713-240-8800

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

BMC SOFTWARE

☐ 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/P.C. _____

Phone _____

MJ

CIRCLE #103 on Reader Service Card ▲

F I G U R E 3

Authorization Table

Authorization Table Format

P S	P S	P S	P S	P S	P S	. . . .
AX = 0	1	2	3	4	5	n ....

Example of AT in Storage:

```
3C000000 0000000 0000000 0000000
```

Where entry for

AX 0 = 00
AX 1 = 11
AX 2 = 11
AX 3 = 00

and all other AX's = 00

F I G U R E 4

XMS Authorization Services Macros

Macro	Description
AXRES	*Reserves an authorization index number*
AXSET	*Set the AX of an address space*
ATSET	*Sets the PT and SSAR authorities for an AX in the AT entry*
AXEXT	*Determines the AX of an address space*
AXFRE	*Relinquishes and returns the AX*

zero). Note that authorization checking is done for *both* supervisor state and problem state programs. By the way, if you can completely understand this explanation by reading through it only once, you may be over qualified for data processing work and may want to investigate getting into a more challenging profession!

Obtaining Authorization

As a default, all address spaces are assigned AX = 0 when they are created. As the table in Figure 3 indicates, AX = 0 is an unauthorized value that prevents the address space from using both PT and SSAR. On the other hand, the default for system address spaces that are using XMS is AX = 1. As a result, system address spaces have full authority to issue PT and SSAR to any active address space. In order to obtain other than a default AX value, the address space must explicitly reserve and then set the AX using the AXRES and AXSET macros which are provided by XMS. Once a unique AX is assigned, any address space to be accessed using this AX must likewise have its AT updated with the correct AX values which is accomplished by using the ATSET macro. ATSET effectively stores the appropriate bit pattern in the AT indicating the user's PT and SSAR authorities associated with that particular AX value.

For a program to use XMS macros, it must be in supervisor state and/or running with a protect key from zero to seven (PKM 0-7). All user address spaces default to use protect key eight, so they are not allowed to use XMS macros. However, in order to make a subsystem's services available to a user, the ATSET macro must be issued from the user's address space. When the user of the subsystem's services is a problem state user, as is the case most of the time, the subsystem must arrange for the macro to be executed on behalf of the user. This is required because in order for the ATSET macro to update the user's AT, the user's address space must be the home address space. Often an SVC is provided by the subsystem that the user can invoke to get the environment initialized properly.

For a complete list of XMS macros that manipulate authorization tables and their contents, refer to Figure 4.

Summary

Since Cross Memory provided such a major addition to the capabilities of the System/370 architecture, the previous protection mechanism had to be updated in order to accommodate this new technology. This included an extension to the protect key facility called the PSW Key Mask. Also, XMS required new authorization checking for the instructions used while in cross memory mode which led to the addition of a new classification of instruction — semi-privileged. Each instruction may be executed by a problem state program. However, it may be suppressed if the user does not qualify under authorization checks performed by the hardware. Finally, in order to provide security and integrity across address spaces using cross memory, XMS assigns authorization indexes which confine cross memory access to a specific community of address spaces.

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.

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

EXCLUSIVE FROM MILTOPE!

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

Cut sheet printers available at 30/60/75 ppm

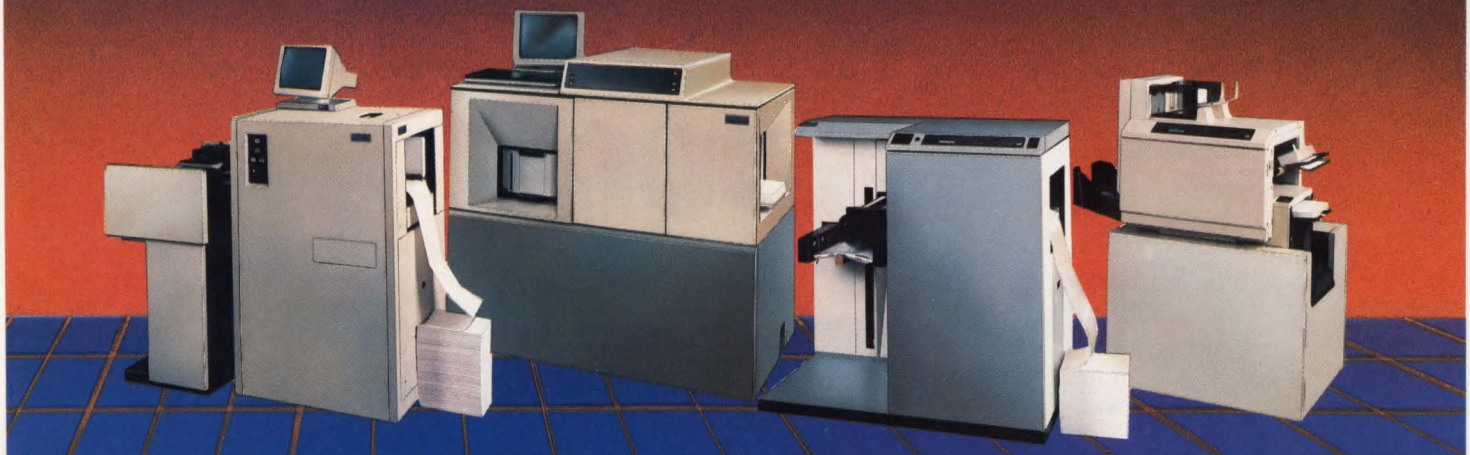

Miltope's FAMILY of Non-Impact 'ion' deposition page printers affords letter quality printing at 'less than a penny' per page.

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

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

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

The choice is YOURS!

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

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

"IMAGINE YOUR IMAGE"

MILTOPE Business Products, Inc.

1770 WALT WHITMAN ROAD • MELVILLE, NY 11747

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

Miltope Printers use Delphax ion deposition engines.

QMS and MAGNUM are registered trademarks of QMS, Inc.

CIRCLE #27 on Reader Service Card ▲

Electronic Report Distribution

By Howard W. Miller

Unattended computer operations have created a market for a new concept with a long history — electronic report distribution.

The popularity of on-line computers and information systems along with the introduction of personal computing has increased the requirement for greater availability of computer-based information. Further, the increasing popularity of electronic mail and the advent of the concept of unattended computer operations have created a market for a new concept with a long history — electronic report distribution.

Over the past decade, virtually all new computers developed have been designed to provide immediate input, update and retrieval of information. Further, the majority of mainframe-based computer application software developed during this same period has been on-line application software. As a result, the business community has embraced immediate response as a mode of doing business and has increasingly demanded more timely computer-based information. This has become a self-perpetuating process. The demand has resulted in the installation of a large base of terminal-type equipment and highly reliable communication networks to tie the base of terminal equipment together.

Personal computing, either on mainframe systems or, as is more frequently the case, on stand-alone personal computers, is another factor that has perpetuated the demand for the greater availability of computer-based information. Personal computing tools have provided the business person with an unprecedented ability to analyze information. However, over the short history of computing, the business community has developed a report mentality: information and reports are thought of as synonymous. Historically, reports were the only format for receiving information and analysis was accomplished by manually transcribing information from a report to an analysis format. Computer information, therefore, became synonymous with reports.

Personal computing is changing this view of information. Business personnel now view information as something electronic that can be moved from one electronic media to another. The whole perception of information is changing.

Electronic mail has also been introduced into many organizations. Electronic mail is reducing the need for hard copy memos or letters, thus speeding the delivery of information and providing a more flexible means of communications. Electronic mail has further demonstrated a disparity between the continued dependence on printed reports for information and the ability to communicate electronically.

Lastly, there is a trend toward unattended computer center operations, operating a computer processing center without any human intervention between the input or update of information and the delivery of the information. A major obstacle to achieving this goal is the continued dependence of most organizations on printed reports as the primary method for distributing information.

These trends have created an unprecedented opportunity to automate the deliv-

ery of computer output. This delivery method is commonly called electronic report distribution. Electronic report distribution is the storage of printed reports on an electronic storage media, usually a disk drive, for some predetermined period of time. The electronically stored information is then retrieved on-line via an existing network of terminal equipment thereby eliminating the need to print reports. Viewing is facilitated by software that permits the manipulation of the electronic report. Upon expiration of its useful life, an electronic document is replaced with the current version, archived or destroyed. In most cases, facilities are provided for low volume, exception printing.

Electronic report distribution is not a new concept; programmers have been using time-sharing tools to view reports, JCL and program code on-line for a long time. Furthermore, some application software systems have provided on-line report viewing for a long time. Until recently, the timing has not been right for the widespread use of this on-line report viewing by the business community. However, the number of installed on-line terminals and personal computers has created a critical mass of equipment that has put a terminal within easy reach of most key employees. This factor, along with the decreasing cost of direct access storage devices, has now made electronic report distribution practical.

It is important to understand that from a strategic point of view electronic report distribution is not a replacement for interactive data query. Electronic report distribution is a bridge to get from printed reports to electronic distribution without a huge reprogramming effort.

Expectations

The requirements for an electronic report distribution system are essentially the same from organization to organization. However, the importance of different attributes will vary significantly based on the types of reports processed. As an example, a user with little or no data processing experience may simply wish to scan a report, selecting information in the same way as viewing a hard copy report. Conversely, an application programmer may desire to manipulate the format of a report to facilitate the validation process and may, therefore, be very concerned with the technical capabilities of the electronic report distribution software.

Some of the requirements common to both a technical and an end-user audi-

ence will include the ability to do the following:

- View reports on-line at an existing terminal
- Select on-line viewing or hard copy printing as the viewing medium
- Limit the access to confidential reports to authorized personnel
- Share one or more reports with one or more users
- Store reports in a compressed and sometimes encrypted database
- Scroll report information left and right when it exceeds the capacity of the terminal
- Produce hard copy reports in whole or in part at local or remote printers
- Produce report archival and restoration
- Install the software with no modifications to operating software, application software or JCL.

Specific operating functions and capabilities vary from one software vendor to another. Further, the importance of a function will vary from one organization to another. Therefore, it is important to develop an inventory of expectations to systematically evaluate the capabilities of different software vendors.

Establishing Requirements

There are numerous electronic report distribution systems on the market. See the accompanying sidebar titled "Electronic Report Distribution Products" for a list of nine of the more common IBM mainframe-based systems. These and other report distribution systems can be differentiated by product capabilities, technical support information, implementation information and miscellaneous information.

Product capabilities may include on-line viewing, remote printing, archival, report forwarding, usability, report distribution, security and product administration.

Technical support information includes product documentation, communications monitor, operating system support and technical support information.

Implementation information includes information about report set-up, JCL changes, hardware requirements, software prerequisites, implementation effort and complexity to change set-up.

Miscellaneous information includes pricing, discounts and maintenance schedules, installed base of users, vendor information and user group information.

In order to differentiate one report distribution system from another, develop a detailed list of expectations or requirements. Developing such a list is the most labor-intense aspect of selecting an electronic report distribution system. However, this effort can be minimized by analyzing vendor literature. The analyst can identify the different features offered by the vendors and these features can be translated into concise statements of requirements. The statements can be grouped into the four classifications identified above. To further facilitate this process, accompanying this article you are provided with a comprehensive list titled "Electronic Report Distribution Requirements Inventory."

After the initial list of requirements is developed and classified, present them to a group composed of user representatives, information technology professionals and

Electronic Report Distribution Requirements Inventory

Product Capabilities

On-line Viewing

Is there a maximum allowable number of pages that can be read on-line?
Can JCL be read on-line?
Is a separate job class required for JCL?
In viewing mode can you search by:
- Constants in the report body?
- Headings?
- Constants by row/column?
- Use of Boolean logic?

Is there a tutorial function?
Is the product capable of viewing all reports (test & production)?
Does the product have the ability to reformat & save reformatted reports?

Remote Printing

Does it provide print driver support for remote printers?
Does it provide mainframe-to-mainframe to PC-to-minicomputer distribution?
Does it provide downloading reports to PCs?
Does it allow users with printers to manage their print queue?

Archival

Is there archival to tape and/or disk?
Is there automatic archiving?
Is archiving done with a batch job?

Report Forwarding

Can the user route a report to another user?
Can the user route a report to another user, masking selected data?
Can the user route selected pages of a report to another user?
Can the user add comments before routing a report to another user?

Users

Is the product capable of being used by non-technical staff?
Is the product menu driven?
Is a tutorial/help function available for each screen?
Is user education adequate for full use of products?

Report Distribution

Are different reports going to the same area grouped together?
Can a report be selectively distributed by searching a row or column for a variable?
Can a report be selectively distributed by searching for changes in heading?
Can the number of copies be altered while the job is running?
Are there separator pages between printed reports?
Is there an itemized list of reports by print order?
Is there an itemized list of reports by recipient?
Can sensitive reports be selected and printed separately?
Are laser printers supported?
Are there overrides for priority printing?
Does software have an SMF exit to collect and store additional data?
Does the software provide a general activity report?
Does the software provide data on late reports?
Does the software provide data on jobs processed or not processed?

Is your PRODUCTION SCHEDULE out of control?

Do RERUNS have YOU on the run?

Perhaps you should consider . . .

CONTROL-M

The Automated Production Control System that puts YOU in control.

- Totally automates Job and Event Scheduling
- Improves system throughput
- Significantly reduces reruns:
  - Ensures jobs run in the proper sequence
  - Submits Jobs ONLY if ALL prerequisites (manual, predecessor jobs, datasets, hardware resources etc.) are met.
  - Automatically updates JCL and Control Cards, eliminating manual errors.
- Provides on-line management visibility and control
- Schedules both batch jobs and Started Tasks
- Automatically opens and closes CICS files
- Automates restart and recovery
- Simulates hardware and software changes
- Puts job documentation on-line
- Produces immediate notification of problems
- Installs in 1 hour WITHOUT ANY SYSTEM HOOKS
- Makes schedule definition quick and easy

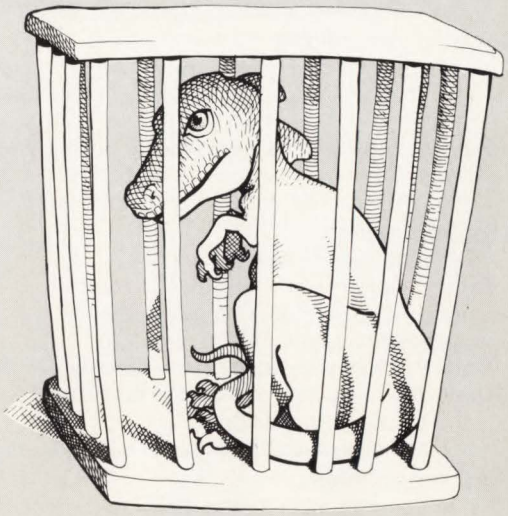

CONTROL-M tames the RERUN monster

MVS, MVS/XA, JES2, JES3

TONE Software Corp.

1735 South Brookhurst, Anaheim, CA 92804-6491
(714) 991-9460 • Telex 4974583 • 1-800-833-8663

Technical Support Information

Support Information

- How long has the product been on the market?
- Is there hot line support?
- Is there 24-hour support, hot line or otherwise?
- Does the product have adequate support and numbers of professional staff?
- Has there been more than one release per year of the software over the last two years?

Product Documentation

- Does it provide on-line documentation for all or selected users?
- Does it provide technical documentation for support staff?
- Is documentation adequate to resolve routine questions without the need for vendor support?
- Are all internals documented?

Technical Issues

- Does the software run with your TP monitor?
- Is there a security interface into your security monitor?
- Are descriptions of user exits available?
- Does it provide logical versus physical viewing of data?
- Does it provide dynamic column and width adjustment of terminal image?
- Does it provide support for your current release of the spool software?
- Are operating system modifications required to install the software?
- Will the software run with your release of the operating system?
- Are storage and access methods compatible?
- Are archival and retrieval methods compatible?
- Can microfiche be created?
- Can it support simultaneous creation of reports and fiche?
- Are all documented internals operational?
- Are special JCL definitions/changes required?
- Are there interfaces with other products?
- Is the software compatible with other products?
- Will it run with other operating systems?
- Are path lengths documented for each function?
- Are there other planned upgrades or enhancements?

Implementation Information

- Are JCL or PROC changes required to install report distribution?
- Are the administrative activities required to set up a report reasonable?
- Is more than one screen required to set up a report?
- Does setup require the assistance of a programmer?
- Is the time required to set up a report reasonable?
- Is it easy to add or change a report?

Miscellaneous Information

Client Information

- Is a list of installed users available?
- Are there installations similar to yours?
- Are technical contacts available for installations similar to yours?
- Is a list of installed users available that have converted to the product?
- Is there an active users group for the product?

Contract Information

- Have you received a copy of the licensing agreement?
- Have you included a product escalation clause for timely resolution of software failures?
- Have you included an acceptance test as a condition of acceptance?
- Have you withheld monies as a condition of acceptance?
- Have you received assurance that the supplier will provide continued support?

Cost Information

- Do you have the product cost?
- Do you have the maintenance cost?
- Do you have the cost of a new release?
- Is a deferred payment plan available?
- Does the license extend to an alternate site, for example for disaster recovery?
- Have you calculated the cost of disk storage (this may vary by product based on the compression formula)?
- Are discounts available (multiple site, education)?

Other Information

- Does the product have the ability to recognize logical punch format output (from the system punch)?
- Can you phase the installation of report viewing and archiving?

management. Make this a brainstorming session and expand the list to include as many functional requirements as possible. Remember, there is no end to the number of items that can be added to a wish list, so be concise and avoid duplication and overlap as much as possible.

The functions of the vendor software are compared to the list of requirements and are numerically evaluated for compliance. Calculations are required. Further, as the process proceeds, requirements are added while others are dropped. It is therefore suggested that the requirements list be maintained on a spreadsheet. In this way, items can be added or dropped, ratings can be calculated and the list can be sorted with little or no manual effort.

The goal of the evaluation is to select the software that best meets the requirements identified by the analyst. When the rating and scoring are complete, compare the score for each electronic report distribution package and select the two packages with the highest rating. Schedule a site visit for each and select one of the two for a 30-day trial period. If the results of the site visit or trial period are not satisfactory, look at the next alternative.

Observations on Selection, Justification and Implementation

Early interest in electronic report distribution systems was oriented toward batch report management and tracking capabilities. Many products were designed to support this objective rather than on-line viewing. Assuming that on-line viewing is the major requirement driving your selection, pay special attention to how flexible or user friendly your finalist is on this aspect. Remember, there are those who designed this facility into the product and those who added it as an after thought. The results are not necessarily the same.

Report preparation and distribution is a time consuming, labor-intense, high cost process that is plagued with errors. Further, the distribution of manual reports is a major security fault point. Distribution of sensitive data by inter-departmental mail or courier is subject to error and embarrassing security violations. An electronic report distribution system is a solution to these problems.

Although cost reduction should not be the driving motivation for the installation of an electronic report distribution system, cost reduction is a great "attention getter." It is an opportunity to raise the

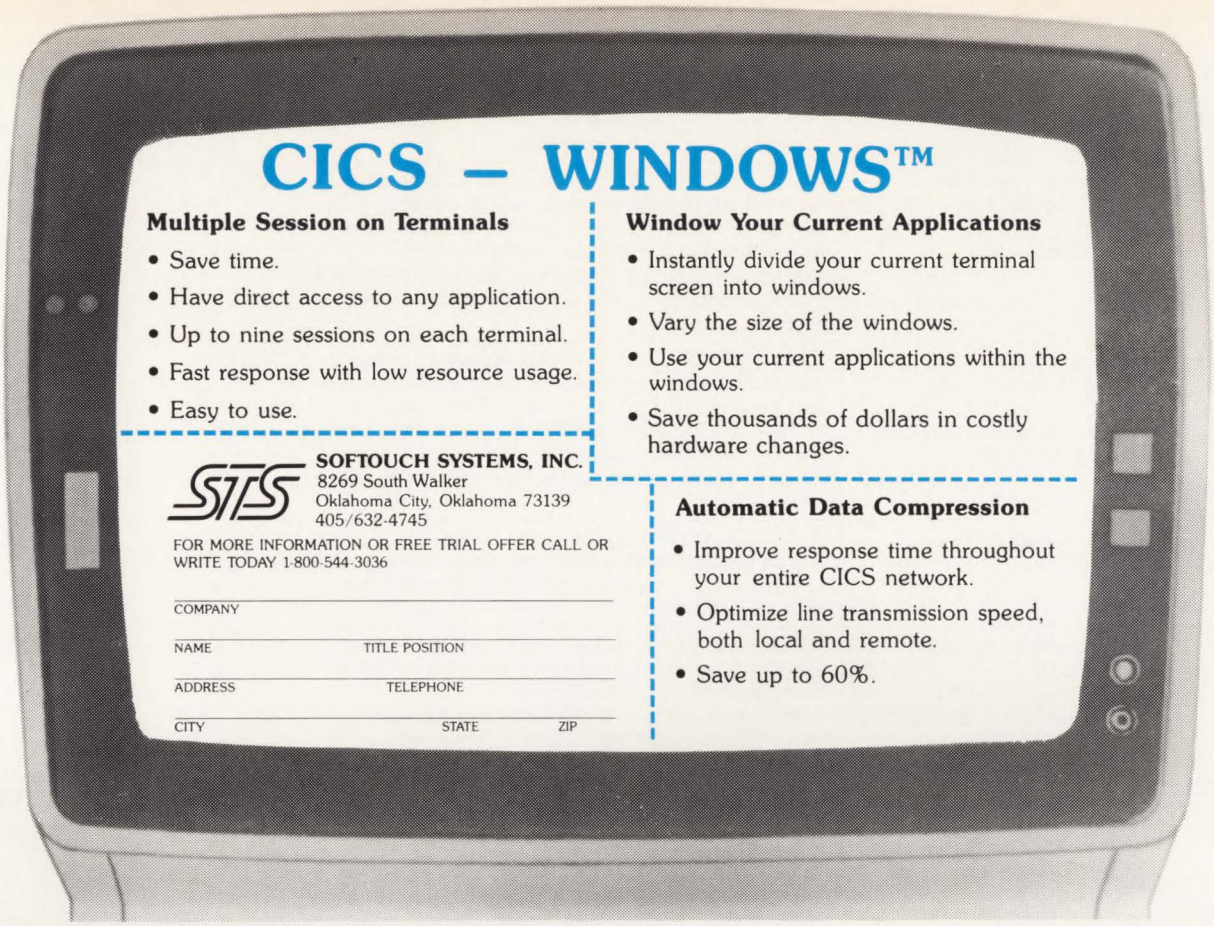

CIRCLE #47 on Reader Service Card ▲

THE REPORT MANAGEMENT SYSTEM

- MVS/SP or XA
- JES2/JES3
- Absolutely No System Modifications
- Communication through TSO, ISPF, ROSCOE, IDMS, IMS, CICS and VTAM

To date, we have helped over 400 organizations address their productivity concerns. If you would like to receive more information or an onsite presentation/demonstration, write or call us.

RSD

THE SOFTWARE COMPANY WITH PRODUCTS FOR TODAY AND TOMORROW

EUROPE: RSD SA 41, rue Marziano
1227 Les Acasias
Geneva-Switzerland
Tel. (22) 43 21 10
Telex: 429 537 RSD CH
Telefax: (22) 43 21 34

USA: RSD AMERICA, INC. 100 Merrick Road
Suite 500 East Building
Rockville Centre-New York
11570
Tel. 800 777-WSF2 (516) 536 8855
Telefax: (516) 763 1020

CIRCLE #42 on Reader Service Card ▲

Electronic Report Distribution Products

Product Name	Vendor	Brief Product Description	Operating Environment	For More Information
CA-DISPATCH	Computer Associates 711 Stewart Ave. Garden City, NY 11530 (800) 645-3003	CA-DISPATCH is a report distribution management system that controls report production from the moment the report is planned to the time the printed report is dropped in the end-user's mail box. It exercises this control by applying user-supplied report management information to the report production process.	MVS	CIRCLE 300
SAR/EXPRESS	Essential Software Inc. 15233 Ventura Blvd., Ste 614 Sherman Oaks, CA 91403 (818) 906-7796	SAR and EXPRESS DELIVERY are two on-line, integrated products designed to form a complete pre-spool report management system. SAR is used to archive, retrieve, view and reprint any type of SYSOUT and EXPRESS DELIVERY is used to address, separate and handle reports for the end user.	MVS	CIRCLE 301
RMDS	IBM Corporation	RMDS (Report Management and Distribution System) is a series of programs providing users with a system to store, protect, view and print, as necessary, system output on demand.	MVS	CONTACT LOCAL REP
INFOPAC	Mobius Management Systems One Sheraton Plaza New Rochelle, NY 10801 (914) 632-7960	INFOPAC is a fully integrated Output Management System designed to sort and collate reports for easy distribution. INFOPAC automates distribution of reports to any output device including: hardcopy printers, on-line terminal networks, personal computers, remote printers and microfiche devices.	MVS, VSE	CIRCLE 302
BROWSEWSF2	RSD America 100 Merrick Rd., Ste 500 Rockville Centre, NY 11570 (800) 777-9732	WSF2 reports are made available to users on-line upon completion of the production job that produces the report. Its features include: data security, instantaneous location commands, report splitting, reformatting, note pads and printing on demand, etc.	MVS	CIRCLE 303
VIEWCOM	StarTech Software Systems 80 Beaver St. New York, NY 10005 (212) 943-9800	VIEWCOM is a report distribution and management system which supports on-line viewing and printing, segmentation and bundling, report archiving, down-load to PC printers and hard disks. Its benefits include a significant reduction in paper and report handling and distribution costs.	MVS, VSE	CIRCLE 304
X/PTR	Systemware Inc. 12770 Coit Rd., Ste 1008 Dallas, TX 75251 (214) 239-0200	X/PTR captures reports from SPOOL, compresses to a pre-defined database, provides full-screen browse (TSO/ISPF or VTAM) and automatic decollation. It also provides automatic bundling of decollated reports, security, automatic archiving and retrieval, and on-line administration.	MVS	CIRCLE 305
TS-RMDS	Tone Software Corp. 1735 S. Brookhurst Anaheim, CA 92804 (800) 833-8663	TS-RMDS provides an effective approach to automated report management and distribution. It provides report decollation, bundling, on-line viewing, archive/retrieval facilities, report and bundle manifests, report accounting, data compression, etc.	MVS	CIRCLE 306
CONTROL-D	Tone Software Corp. 1735 S. Brookhurst Anaheim, CA 92804 (800) 833-8663	CONTROL-D features scheduling of events, user notification when reports are ready for viewing, laser printer compatibility, printer workload balancing, etc. It can be totally integrated with CONTROL-M, an automated job scheduling product to provide total control of workflow throughout the system.	MVS	CIRCLE 307

interest level of users and management alike. An electronic report distribution system is an opportunity to address the *real* cost associated with hard copy report distribution:

- Computer hardware processing cost
- Print hardware and supply cost
- Mail cost either inter-departmental, external or courier
- Delay in the distribution of information and the corresponding lost opportunity cost
- Training and error cost as a result of personnel turnover
- Management cost associated with explaining why reports are lost or delayed
- The cost associated with a loss of confidence in data center directions as a result of lost or delayed reports
- The cost of reruns or report recovery.

Electronic report distribution systems are an opportunity to improve administrative procedure; they reduce dependence on data centers. Some distribution systems advertise complete end-user administration capabilities without data center personnel intervention. Most products provide on-line viewing, exception printing and volume printing at remote locations that eliminate dependence on data center personnel for printing and distribution.

Some user guidelines are appropriate for implementation of an electronic report distribution system.

The implementation process is an opportunity to assess the need for a report, its frequency and distribution. It may be possible to eliminate or reduce the frequency or distribution of a report.

Set up guidelines that define excep-

tions: retention default; archival default; data center hard copy printing volumes; remote hard copy printing volumes; and service level standards.

Periodically report exceptions to guidelines to the user community and management.

Use departmental coordinators. Provide the initial training to the departmental coordinators. The coordinators provide direct training and consulting to their department.

Initiate the implementation of the electronic report distribution system in the data center or even better in the whole MIS group. Eliminate the printing of *all* reports in the data center. Establish a do-as-I-do attitude, not a do-as-I-say attitude. This will go a long way to developing a positive user attitude.

Communicate the tangible and intangible benefits of an electronic report distribution system to management. Make management understand that it is integral to the objectives of unattended operation and that it is a bridge to on-line information access. Do not let management get caught up on outdated notions about the cost of DASD. Identify that the cost of storage is going down and the cost of manpower is going up.

When implementing, start with the high volume weekly and monthly reports and with the most receptive users of the technology. You are looking for the areas that have the highest return and will generate the most positive impact.

Summary

Electronic report distribution is an excellent opportunity for any organization with a large installed base of on-line terminals and a large volume of printed reports to improve the productivity of its users. The ability to eliminate the manual effort and errors associated with the manual production and distribution of reports, to reduce expenses through the elimination of paper and printer related costs and to expedite the distribution of computer generated information is enticing.

Installing electronic report distribution software is not particularly difficult. Unlike other software selection projects, organizational requirements for electronic report distribution do not vary much from organization to organization. Further, software variations tend to be variations in approach rather than in real capabilities. The major inhibitor to installing electronic report distribution software tends to be the lack of awareness on the part of data center management of the potential contained in electronically distributing printed media.

Lastly, there is a movement toward unattended data center operations. A major obstacle to achieving this goal is the continued dependence of most organizations on printed reports as the primary method for distributing information. Electronic report distribution eliminates this obstacle. It eliminates lost and misplaced reports, reduces job rerun and accelerates the availability of computer-generated information. In addition, it eliminates manual intervention by data center personnel while improving security.

ABOUT THE AUTHOR

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

Now, for your VM environment:

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

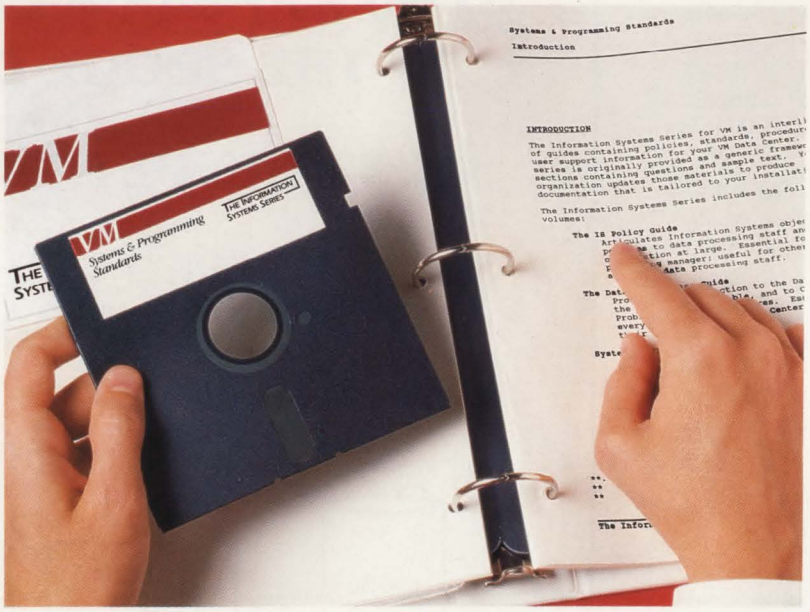

Here's what you've been waiting for! Now, with the ***Information Systems Series***™ documentation framework, you get everything you need to quickly develop custom policies, standards, procedures, and user support documentation for your VM installation.

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

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

Here's what you'll receive.

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

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

Call for more information.

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

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

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

CIRCLE #69 on Reader Service Card ▲

DOCUMENTATION?

WHAT DOCUMENTATION?

By Judith L. Glick-Smith

Upon seeing a dog walk on its hind legs for the first time, Samuel Johnson remarked, "It's not done well, but you're surprised to see it done at all." The same can be said of the documentation in most data processing shops.

How many times have you heard or said, "Just start coding. We'll do the documentation when we finish implementing the system." Then one year later, two months after the system is installed, "Don't worry about the documentation. We need to hurry and get this next phase implemented. We'll document later." Somehow "later" never comes.

Change requests begin to pile up. The people who originally worked on the system have migrated to other projects or other companies. Staff, unfamiliar with the original project, are charged with maintaining it. They spend days wading through poorly documented code and little, if any, documentation.

Does this scenario sound familiar?

Documentation has become an industry joke. Most shops understand the necessity of user documentation. Users will not use the system if they do not have documentation. However, systems and programming documentation, programmers' reference materials and standards and procedures are an afterthought and often do not get written at all. Finally, when documentation, including users' procedures, is written, more often than not it is not maintained on a regular basis.

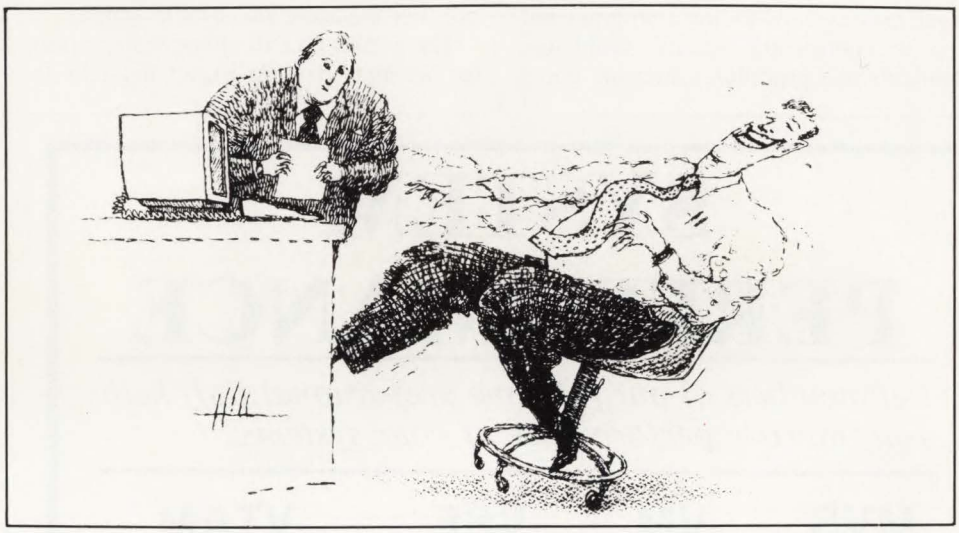

Dominos

The biggest complaint I hear from programmers and systems analysts is that their shops have no current documentation. Stress levels are constantly high. Systems development can be compared to setting up dominos and watching them fall.

New development teams must work from sketchy specifications and little, if any, contact with those people who will ultimately use the system. Decisions often do not get written down and are, therefore, forgotten. The programmers themselves are on a deadline and do not have time to document their code. Time for documentation is rarely set aside in the development plan.

When it is time to implement the system, someone realizes the users need something to tell them how to use it. One of the programmers is assigned to write the user procedures. He does a screen-print of all the screens in the system and takes them to the manager of the user department (who probably understands what is to be keyed in all the fields). Then the users come up with their own procedures. This is not good PR for the computer services department!

The users begin using the system and the real fun begins. Because their needs were not understood from the beginning of the development process, major changes need to be made to the new system. But what has happened to those who

implemented the system? Burnout from long hours, new projects and backed-up maintenance from other systems have scattered your original development team. You must assign personnel unfamiliar with the system to maintenance.

In the meantime, because operations procedures were slim, the operators keep unintentionally messing up the nightly batch jobs. Every morning is a nightmare trying to maintain the integrity of the files. The first half of the day is spent putting out fires; that leaves little time for concentration on change requests.

Everyone in data processing understands this pattern and yet it is repeated over and over again when the situation can be so easily remedied.

Quantifying the Problem

Current systems and programming documentation keeps your programmers productive and saves on maintenance costs. Programmers' reference materials inform staff of shop dependent information such as job classes, how to link a program and who to contact for specific problems. Standards and procedures alleviate guess work as to how management wants to run the shop. All three types of documentation serve as training materials, keeping interruptions of seasoned employees to a minimum.

Intuitively you know that if your documentation has been written and is current, it is saving you money. But how do you measure the benefit? And what if the documentation you have is not current or, heaven forbid, non-existent? How do you justify the cost of writing and maintaining it?

Just as in measuring the benefits of computer systems, measuring the benefit of documentation is subjective. However, it is easy to measure the lost time of not having documentation as can be seen from the following scenario.

Only two people are expert on your accounts payable system. One is on vacation. During this time, the system goes down at three a.m. during month-end processing. There are no operations procedures. The operator calls the person on call (not the accounts payable expert).

The person on call attempts to contact the accounts payable expert but fails to reach him. The on-call person must drive into work to try to determine how to handle the problem.

The quantitative analysis for this situation is quite simple. Multiply the hours lost by the on-call person, the operator and possibly the users in the morning times the hourly rate of those employees. This is the total time lost because of inadequate documentation.

It is difficult to quantify subjective costs, such as the effect of high stress levels for all involved and the effect of resulting bad feelings between the user department and computer services. Conversely, it is just as difficult to quantify the benefits of having documentation.

The Solution

Good documentation does not prevent the system from crashing. It serves as a problem solving tool for your staff. By performing several tasks, you can help your shop move away from old habits and toward a more productive, pleasant place to work.

Set documentation goals. Also, come up with a plan to update current documentation and write documentation where it does not now exist. Include in the plan a commitment to provide for documentation development at every stage of the system development and maintenance process.

Sell the plan to upper management. It will not work without management commitment. Then, begin documenting at the beginning of the system development process and maintain the documentation on a regular basis. Lastly, institute controls to be sure the documentation is being written and maintained.

Set Goals

When setting documentation goals, be realistic. You cannot make current all the documentation for 10 systems in six-months time and still continue to run your shop efficiently. Set a goal for two-to-five years, depending on the size of your shop, to have all the documentation for all the systems current. Then ask yourself, "What must we do to make that happen?" In other words, treat this project like any other system development project.

Come Up with a Plan

Identify someone to head up the documentation project. Every shop I have ever been in has at least one person who enjoys doing documentation and does it well. That is the person to lead this project. It

is not important that this person has never led a project before — what is important is this person likes to do documentation.

This person should have no other tasks. Give change requests and system responsibilities to someone else. You say you're low on resources? Ask yourself, "How committed am I to reaching my goal of a fully documented shop?" By the way, do not change this person's job title. This is a project leader analyst position — not a technical writer position.

Assign the new project leader four tasks. Identify the areas that need documentation. Then, analyze the cost of having poor documentation in those areas.

Estimate the time it will take to do the documentation, then double the estimate. This is done because so often we fail to remember publication time. In the case of manuals, publication time includes proofing, printing and distribution. In the case of on-line documentation, publication time includes choosing and implementing the on-line package.

Based on time estimates, come up with a cost estimate. Do not forget the cost of binders and printing if the documentation is to be in manual form. If this is to be a purchased package for on-line documentation, include the cost of the package.

Based on your project leader's findings, decide how the project should be staffed. Depending on your time frame and the availability of resources, you may want to contract out some of the work. An employee should head the project, so that after the contractors are gone there will still be a commitment to maintaining the documentation.

Hint — do *not* use Journalism or English majors to write your systems and programming documentation unless they have at least two years systems and programming experience. I have found that programmers and systems analysts who made A's in English in school make the best documentation specialists.

Sell Your Plan

The most important step in correcting the problem of poor documentation is the same as for system development: obtain management commitment. Use the information gathered by your project leader and the plan you have developed to present your case. Compare the high costs of not doing the documentation with the cost of creating and maintaining the documentation.

When the up-front cost is high, you must convince management to think about long-term costs savings after the documentation is in place. Using actual examples, demonstrate how current documentation will pay for itself in a relatively short period of time.

A long-term plan such as this cannot work without the commitment of management. In order to reach your goal, resources cannot be reassigned at random.

Document Systems Development

Every shop should hire or assign someone to be documentation specialist whose primary task is to develop and maintain documentation. Use your documentation specialist in *all* system design and development projects at the beginning of the project. This person can write detailed specifications based on information he gathers during design and planning meetings. Implementation of systems is much easier when all the players know exactly what to expect.

Including your documentation specialist in the beginning phases of system development also makes the job easier when it comes to writing systems and programming and operations documentation. From his initial specifications, test plans and user training and user reference materials are more easily developed.

Maintain the Documentation

Once the documentation is written, it must be maintained to retain its credibility. As change requests get processed or standards and procedures change, the documentation specialist updates the documentation. Having one person or a specific team of persons assigned to maintain documentation ensures that updates will be consistent in form and content.

Again, use systems- and- programming-oriented people. They tend to know what questions to ask and to recognize when there are holes in information given to them.

Institute Controls

Just as with any other system, your documentation needs to be documented. It, too, must have controls and procedures for ensuring that the work is being done. Management should periodically review the work being done by the documentation team. Fortunately, documentation is a high-visibility item. If you know systems are being changed and you are not receiving updates to your systems and programming documentation, your documentation team is not doing its job.

Conclusion

Contrary to popular belief, good documentation is not the impossible dream. In fact with good planning, diligent implementation and maintenance, it is not only possible — you will believe it to be a godsend. Documentation increases productivity, reduces maintenance costs, makes users and systems employees happy and literally lets you sleep better at night.

Don't you love a win-win-win situation?

ABOUT THE AUTHOR

Judith L. Glick-Smith is the owner of Integrated Documentation. She has a BBA in accounting with a minor in information systems from Georgia State University and has been in data processing for eight years. The first four years she spent programming in IBM mainframe shops where she had to maintain too many undocumented systems. Her company, formed in 1983, specializes in writing internal documentation for data processing shops. Integrated Documentation, 2995 LBJ Freeway, Suite 200, Dallas, Texas 75234, (214) 888-1003.

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.

Caching Controllers Improve Computer System Performance

Caching combines the speed of electronic storage with the capacity of magnetic media by utilizing intelligent storage controllers.

By Dennis Corbin

As businesses continue to rely more heavily on technology to achieve success, computer system investments quickly become a costly competitive resource. In order to maximize system performance, most companies are now investigating a variety of enhancement alternatives to extend the performance and life of existing computers.

Recent changes in computing technology and use have made accessing, not processing, data the most critical concern for performance-minded MIS installations. Throughout the last decade a trend occurred that enabled vendors to double processor speed every few years. This no longer holds true. Today, system components are the key to increasing the availability of data to the CPU, not new generations of CPUs. Of these alternatives, caching has proved to be one of the most successful.

An Effective Alternative

Early computer applications were primarily scientific in nature and involved complex calculations that required great CPU resources. As datasets grew larger, users required more powerful processors to handle ever-increasing tasks. More processing power meant faster response time and quicker performance for the user.

As computers became more prevalent in commercial business, the role of the computer system changed. Where computers were previously required to perform sophisticated calculations on the same dataset, the new breed of users needed to perform relatively simple operations on a very large amount of data, such as updating information on a worldwide customer list. These users were primarily concerned with accessing information, not processing it. The fast, effective transfer of data became even more critical when improvements in magnetic storage technology allowed larger and larger databases to be stored on-line.

In this environment, upgrading to the next-biggest processor was not always the solution to response time problems. Only by exploring all the components in a computer system, from disk and tape drives to terminals, could continual improvements in performance be made.

Performance imbalances were linked to the tremendous speed differences between electronic storage devices and mechanical devices such as tape and disk. By virtue of design, electronic storage is thousands of times faster than magnetic media (tape and disk). However, the increase in speed is compromised by capacity limitations. More sophisticated computer technology has evolved to bridge this technology gap in the form of cache storage.

Speed and Capacity

Cache is high speed, electronic storage which retrieves data from other slower devices such as main memory or disk drives. By combining features of hardware and software, cache or storage controllers anticipate the data to be requested next by the CPU and transfer it from disk to much faster electronic storage. As a result, cache reduces average disk transaction time and increases overall throughput. This eliminates the bottleneck traditionally associated with the interface between the CPU and disk drive.

Ideally, cache holds a system's most frequently accessed data. The storage controller contains high-speed electronic subsystem storage which has two main areas: cache and directory. The cache is used to store pages of data for quick processor reference and the directory keeps track of pages stored in cache. The cache controller manages these two areas to form a high-performance paging and swapping subsystem.

Caching addresses two of the fundamental mechanical delays associated with disk drive operations: seek time and latency. When a request is made by the CPU for data stored in DASD, the time it takes for the read/write head to position itself at the correct cylinder is referred to as a disk seek. Latency or rotational delay is the time required for the appropriate sector to rotate under the head. Average seek times can range from 16 to 80 msec and latency averages eight msec.

An effective caching system improves DASD reliability by reducing the number of seeks performed by the drive unit. The result is reduced wear and tear of mechanical drive components and faster transfer of data to the processor. This improvement in response time provides commensurate increases in system performance to maximize user productivity.

Caching logic incorporates locality and measurement of "least recently used" (LRU) algorithms to reduce system response time. Locality operates on the assumption that data once used is likely to be re-used and if files are at all sequential, the adjacent data is likely to be requested next by the CPU. When the CPU request for data is found in cache, it is referred to as a cache "hit." A cache "miss" takes place when the request is not found in cache and a disk access must follow. Caching is most efficient when it realizes a high percentage of "hits," since the slower mechanical disk access of DASD is eliminated.

As more reads are performed, the information in cache is updated with the "least recently used" data being overwritten with the new data. The cache is managed by internal algorithms that decide which information to keep in cache or replace. This approach keeps the percentage of cache hits as high as possible. At the same time, it makes it possible to handle a much larger set of data than would fit in cache.

Caching Trend More Prevalent

Only a few years ago, the cost per megabyte was so high (about $20,000 per megabyte) that adding enough cache to improve system performance was hardly a cost-effective solution. It was more economical for users to invest in additional low-end controllers, rather than add cache to the current controllers. Another drawback that prevented cache from becoming widely accepted was the cost of tuning required to manage files on disk so they could be sequentially accessed for caching purposes.

Today the cost per megabyte has decreased to the extent that adding enough cache to realize an improvement in performance is not as cost-prohibitive as it used to be. Since cache size can affect subsystem performance, the larger the cache, the more data it can hold. Having more data in the cache increases the probability of finding the requested data there which means the tuning associated with file management becomes less important.

Disk controllers with caching capabilities became part of IBM's mass storage plans in the early 1980s. IBM's 3880 model 21 and 23 controllers, introduced in 1985, both contain caching features, as does the 3990 Model 3 due out at the end of 1988.

When Caching Is Effective

A system temporarily plagued by bottlenecks and contention for shared data will be better balanced as improvements in DASD response time allow people and components to work more efficiently and productively. Caching controllers increase application throughput, provide faster response to present users and make it possible to support additional users.

Systems with very high numbers of transactions, particularly within the airline, financial, hotel and data service industries, are very sensitive to service times and use of DASD storage. Within these industries, certain applications will experience significant performance benefits from caching controllers. Some of these applications include: Informational Management Systems (IMS), Customer Information Control System (CICS), Time Share Option (TSO) and Computer Aided Design-Computer Aided Manufacturing (CAD-CAM).

Most successful businesses have made major investments in their data processing resources and are constantly investigating ways to increase system and user productivity. Caching controllers efficiently manage the transfer of data from DASD to reduce response time to the end user. In a competitive data processing environment, increasing user productivity and system performance saves time and, therefore, money.

ABOUT THE AUTHOR

Dennis Corbin is manager of technical sales support at EMC Corporation, Hopkinton, MA. He has more than 20 years experience in the computer industry including work with IBM and National Advanced Systems. Currently, he works with EMC's mainframe products group.

A Powerful Message For Cache Conscious 3880 Users.

Introducing EMC's PowerCache-16 Cache Storage Upgrades.

One of the reasons you invested in an IBM 3880 Model 21 or 23 Storage Director was its caching features.

But with the high cost of IBM cache upgrades, configuring enough capacity to make caching an effective way to improve DASD subsystem performance has been difficult to justify to yourself or to your management.

Finally 3880 Users Can Increase Capacity And Decrease Expenses.

Well all that has changed. Because now there's the PowerCache-16 cache storage upgrade from EMC Corporation. Priced 33 percent less than IBM's comparable upgrade, EMC's PowerCache-16 gives you an affordable way to increase the cache capacity of your Model 21 or 23 Storage Director.

Price Comparison

Upgrade Capacity	EMC List	IBM List	Savings with EMC
16MB	$ 42,000	$ 56,000	$14,000
32MB	84,000	112,000	28,000
48MB	126,000	168,000	42,000

Your Cost Savings Are Only The Beginning.

EMC also uses megabit RAM components and a newer design for higher reliability. Then we put every upgrade through a 100-hour test and burn-in procedure that includes qualification in one of our 3880's prior to shipment to you. Finally, we back our PowerCache-16 upgrades with a no-cost lifetime warranty and choice of service plans.

And, like the thousands of IBM users who already rely on our electronic storage enhancement products for their computer systems, you'll find that use of our upgrades has no effect on the level of service you receive from IBM. That's IBM's published policy.

Cache In On System Performance.

To learn more about the first affordable upgrades for your 3880 Model 21 or 23 Storage Director simply call EMC at 1-800-222-EMC2. (In Mass. 508-435-1000), or write EMC Corporation, 171 South Street, Hopkinton, MA 01748. We'll answer all your questions and put you in touch with the EMC office nearest you.

Call Today: 1-800-222-EMC2

In Mass. call 508-435-1000

EMC²

The System Enhancement Company.

Printed in U.S.A. 6/88
L88-044

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 five 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

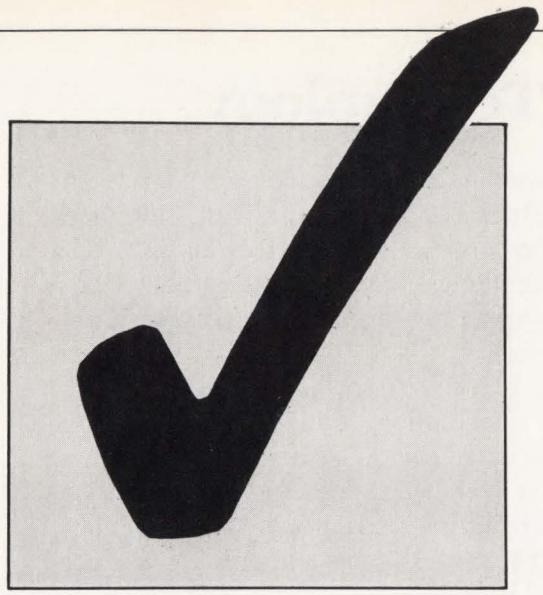

Checklist for Disaster Recovery Planning and Testing

By Kern Chang

A tremendous storm hit a large metropolitan area. The heavy downpour played havoc with the functioning of several local data centers.

A fire raged through a business and shopping plaza in a major metropolitan area. This disabled the data processing service of one of the area's largest conglomerates impacting more than 32,500 people employed in its 15 divisions.

These companies were the fortunate ones. They had disaster recovery plans in place and backup sites selected. Today, many DP shops do not even recognize that their entire data processing operation can be wiped out by something as simple as a water leak (see sidebar).

Stories such as these demonstrate that disaster recovery planning is a vital part of contingency planning. Disaster recovery enables data centers throughout the world to react to the many kinds of outages that occur regularly in the production data center environment.

Contingency planning has recently been heavily promoted to data processing executives by some major vendors. Awareness of this concept is no longer an issue. The biggest issue facing DP executives today is to define the scope of this effort and to start implementing the contingency planning process.

The intent of this article is to review the concepts and potential implementation plans for disaster recovery planning and testing.

Getting Started

Several key requirements for the implementation of a disaster recovery plan include a mandate from senior management and a commitment from senior management for funding.

A third requirement is a commitment from senior DP management to provide:

- Technical staff to head the project and provide direction
- Development staff to provide important input and recovery documentation
- Operations staff to provide support for scheduled backups and off-site vault management
- Communications staff to design and implement network backup/recovery strategies
- Users to prioritize applications.

Last is services such as software/hardware vendors, office supplies, office furniture suppliers, legal consultation, accounting, transportation and so on.

To fully develop and implement a disaster recovery plan may take several years. After the plan is developed and in order to make it a real *working* plan, regularly scheduled testing and updating will demand more resources from the organization. This includes semi-annual testing in addition to the requisite testing whenever a change is implemented and key walkthroughs of the plan. All of this is an excellent investment in terms of company survival.

Obviously, the kind of involvement required from every facet of an organization must be endorsed from the highest levels.

Backup Action Plan

Before a recovery plan can be fully developed, a backup action plan must be in place. The plan should include:

- A complete inventory
- List of the vital records location
- Master control item list
- Vital record update procedure
- General overview diagram and item number detail
- Disaster recovery plan development/team diagram
- Master checklist
- Offsite storage access change procedure
- Offsite storage for hot site/shell operations
- Test plan document
- Disaster recovery documentation
- Disaster recovery action plans
- Active criteria checklist and procedure

Cellular Technology

Cellular technology can eliminate the communication problems often experienced during a disaster. There is nothing more frustrating than attempting to execute your emergency telephone tree during a declared disaster and losing all communications with your disaster team members. It is only logical to expect regional phone outages during wind, ice and fire-related incidents.

The American Cellular Phone Industry Association predicts that by the end of the year more than one million of us will have mobile phones. This technology can no longer be considered a luxury.

True, cost justifying a dozen cellular phones for your business resumption plan with an average cost of $1,295, a fixed monthly service fee of $35 and a usage charge of 35 cents a minute can subject you to an audit review.

There are, however, alternatives. One is to identify and maintain a list of individuals in your organization who have car phones. Also, negotiate a rental arrangement for an adequate number of phones in the event of a declared disaster. Another is to purchase a minimum number to establish communications between the command center and the recovery site. Last is to rent automobiles with car phones. Currently Avis, Budget, Hertz and National offer this service in most major cities.

The rainstorms and switching station fire in Chicago saw one major hot site vendor scrambling to locate enough mobile phones to keep recovery operations running smoothly for their clients.

Features to consider when selecting your cellular phone are: memory dialing for executing telephone trees, minimum one-hour portable battery packs and speaker phone.

Beepers, citizens band and ham radios are restricted by signal strength, quality of transmission and eavesdropping. Clearly, cellular technology offers the most secure and reliable form of communications for the money.

By Tari Schreider, president.
Contingency Planning Research,
Greenwood Landing, N.Y.

- Operating system(s) standards
- Disaster recovery support manual
- Project management guide
- Current building access list
- External vendor contact list
- Corporate/DP organization chart
- Insurance procedure
- News release procedure
- Internal contact list
- Critical/non-critical system list
- Audit procedure
- Accounting procedure
- Control center setup procedure
- Recovery team organization chart
- Team leaders' action plans
- DP directory
- Application analysis worksheet
- Senior management support memo
- Support services memo
- Special equipment memo
- Form letter to vendors
- Security for software
- Fixed assets inventory
- DP diskette backups
- User diskette backups
- IBM software directory
- Lease agreements/contracts
- Building floor plan
- Security manual/phone directory
- Office supplies form
- Applications systems inventory and summary
- Job diagrams and updates
- Vendor software manuals and user manuals
- Computer room power down procedure
- Fire suppression procedure
- TMS/VTOC daily listings
- Power cable scheme
- Hardware specifications
- Data control information
- Daily run books
- Weekly run books
- Monthly/yearly run books
- Vacation/lost time distribution
- Hardware detail list
- IPL procedure
- Systems/Applications backup/restore procedures
- Systems software by categories
- Communications locations and critical circuits
- Peripheral failure procedures
- Network recovery procedure
- Power failure recovery procedure
- User liaison team manual
- Administrative support manual
- Hardware acquisition team manual
- Facility reconstruction team manual
- Senior management damage assessment team manual
- Offsite storage team manual
- Hot site floor system plan
- Data center emergency procedure
- Evacuation plan

The plan should also address regular backup and update procedures to keep the above listed information current.

Recovery Plan

The following list is a twelve-step recovery plan:

- Disaster occurrence
- Activation of emergency procedure
- Notification
- Recovery team assembly
- Damage/impact assessment
- Recovery decision
- Activation of the recovery procedure
- Transition to hot site
- Emergency restoration
- Resource restoration
- Transition to shell
- Restoration of normal services.

Disaster occurrence contains instructions regarding identification of a specific type of event (such as a fire) and what steps should be taken. For example: sound fire alarm, evacuate premises, contact fire department and contact building security.

Emergency procedure activation indicates which emergency procedure is to be activated immediately after the disaster. For exmaple: total evacuation, partial evacuation, start fire fighting, emergency system shutdown, move sensitive resources and placement of protective cover over equipments.

See Disaster Recovery *page 114*

VENDOR PROFILE

Users of IBM mainframes sometimes would like to know more about third-party vendors. Vendor Profile is a regular forum whereby different vendors are given the opportunity to introduce themselves and their products to MAINFRAME JOURNAL readers.

Bennett Software, Inc.

Attention to development is the key to success.

Founded in 1984 primarily as a consulting firm, the J. William Bennett Co., Inc. provided specialized IBM MVS system programming services to several Big Eight accounting and consulting organizations, MIS consulting groups such as CGA and on various government projects.

During his nine-year tenure with GTE Data Services, Bennett became interested in designing a TSO-based batch job scheduling system for IBM's MVS operating system. Bits and pieces were developed but the project received little support from GTEDS home office.

In 1986, Bennett began re-evaluating his GTE scheduling concept. Completely redesigned for IBM's TSO/ISPF interface, an initial version of the JOBTRAC scheduling system was evaluated and subsequently purchased by several of Bennett's Houston-based consulting accounts.

The first interstate sale did not occur until May 15th, 1987. In the year that followed, JOBTRAC captured 10 percent of the market for new MVS job scheduling systems and generated more than $1 million in revenue. The company grew to 16 employees, four sales offices and current projections indicate a probable $4 million second year.

Product technology alone will not produce results like these, especially considering there was no funding available for marketing and virtually no advertising.

A definite factor in Bennett's surprising success was the acquisition wars waged during 1987. Cambridge System Group fell to UCCEL and UCCEL, in turn, to Computer Associates gathering the three most visible job scheduling product competitors under the CA umbrella. Customers and prospects concerned about the future of acquired products began looking at alternative products they would not have normally considered.

Dedication to Customer Service

A dedication to customer service from marketing to technical support has given Bennett Software a reputation beyond its years. The company used a strategy of telemarketing-lead generation and a unique product presentation technique using a dial-up session to a live installed customer's production control system. Many prospects became impressed with the power and flexibility of this newcomer.

With the intitial success of the introductory JOBTRAC product, Bennett turned to a strategy for getting a better grip on the specialized field of MVS operations management.

Close attention was paid to each and every customer as they installed and implemented the system. A detailed enhancement and product requirement report was required of each marketing and support employee.

A restart and rerun system was developed to a customer standard that required no JCL or Parm alterations and had to provide ISPF restart control. RUNTRAC was fully integrated into JOBTRAC in May, 1988.

The REPTRAC report distribution system is scheduled for release in the 3rd quarter of 1988, and, once again, is being designed to customer specifications. A system that requires no JCL or Job alterations, REPTRAC will allow the extraction and/or combination of reports or report segments into user-defined sets. All functions are accomplished in real-time with complete archival and recall features.

JOBTRAC/V2 is set for release in the fourth quarter of 1988 and is billed as the most powerful single integrated system available for MVS production control automation.

Standard features in JOBTRAC/V2 include: on-line Sysout management and archival; job scheduling capabilities for more than 10,000 daily jobs; event, message and job dependency scheduling; operator command scheduling; unattended MVS console operations; unattended IMS master console operations; on-line, real-time, JCL edit syntax checker; and more than 60 other integrated services all under authorized ISPF control.

Optionally, JOBTRAC/V2 can come equipped with the NJE Internetwork Scheduling Option (NJE/ISO). ISO allows a single JES NJE network scheduling image. An entire worldwide NJE network of data centers can be scheduled as a single resource. Complete with internetwork job and event dependencies, ISO has no peers in the marketplace.

Specialist in Production Control

Bennett wants to be considered a specialist in production control systems development, but does not want to become too exclusive. Wishing to avoid the lure of providing only powerful systems to only large customers, Bennett announced JOBTRAC V1Basic.

The V1Basic system is a full-featured scheduling system for 4300 type installations and other mid-range users that require ISPF based production control over schedules with several hundred daily jobs. Without all the whistles and bells and weighing in at less than $25,000, V1Basic has all the scheduling power of its big brother without the additional automation functions that push the full blown JOBTRAC/V2 to more than $65,000. Add another $10,000 per site for NJE/ISO.

Realistic Pricing

Realistic pricing is another earmark of Bennett products. Considering the cost of JOBTRAC/V2 with NJE/ISO, RUNTRAC and REPTRAC, as an integrated system, they still come in at less than a third the price of Computer Associates' CA-Unicenter.

In 1989, Bennett Software will introduce Version 1 of TAPETRAC, a new tape and cartridge management system. In 1990, Bennett will release a multi-system, real-time, resource manager interface for all Bennett Products. An MVS subsystem, code named SYSTRAC, is in its initial design phase and promises to be the final act in the unattended computer operations overture.

Attention to Development

As can be seen, Bennett believes attention to development is the key to success in operations management software. The key to keeping development in tune with demand is close attention to customer and prospective customer "wish lists."

Utilizing the marketing and support department's diversified staff, a set of logistical trend reports is processed monthly. Enhancements to current products are categorized by the development staff as to their complexity and ease of integration. The marketing and support managers prioritize the same enhancements by prospective customer demand.

Enhancements and extended features are frequently available in quarterly or semiannual refresh releases at no charge to current customers.

Since the first software sale, all Bennett products have been accompanied by a 100 percent money back guarantee and not one customer has ever requested a refund. Bennett endeavors to ensure that none ever will.

Was Your CIO An Accountant?

MIS executives must stem the tide in current CIO selections.

By R. Douglas Swords

In any business endeavor, whether it is healthcare or manufacturing, there is an attempt to offer goods or services to a market at a price that is attractive and at the same time at a cost that can generate enough of a margin to keep the organization vital. The identification of controllable cost factors has been accomplished by cost accountants, the identification of markets by marketing teams, the identification of new products or services by product developers and underneath it all support by information systems of some description.

Through all phases of business, MIS has had some presence although that presence has not always been popular. As the capabilities of information processing technology have improved, so has the penetration of MIS into the arteries of business. This penetration has been so deep in some areas, for example the stock market, that problems such as the market correction in October of 1987 have been attributed to the systems themselves. This may not be an entirely fair conclusion, but it does indicate how significant the processing of information has become to business in the past 15 years.

The diversity of operations for any business coupled with decentralized operations (headquarters here and branches everywhere) has contributed to the increasing complexity of applications and the networks to connect them. Business planners have had to consider the capabilities of computers in their expansion planning. The cost of systems becomes an important factor in the cost of doing business. The value of information has exploded onto the forefront as one of the major business issues of the next decade and century.

Strategic Weapon

The phrase ''information is a strategic weapon'' has become popular. Information has always been a strategic weapon. The problem was in identifying what portion of overall information was strategic for a particular business. The cost of doing business five years ago might have been important then, but its value is limited now except as an indicator of inflation or perhaps whether or not new efficiencies in production have been found. Information pertaining to *what* that business was doing then is critical and clearly has strategic implications: when compared to current information it will indicate the direction of the business.

An example might be the cost of the production of spoons. Spoons were the only product five years ago, but now we also make forks and knives and we make them in plastic, silver and pewter. We have diversified our product line, we buy more raw materials, we have greater economies of scale. This is an over-simplified example, but it points out that within the set of all information only a subset is strategic. The cost of making spoons was clearly far more critical when we only made spoons. Now, the fact that we have diversified is more important and the reasons for diversification become critical. Those reasons form the basis for strategic planning. Did the factor that drove our diversification come from inside or outside the organization? Was it a market reaction? When did we change?

CIO

The reason that this issue is important to MIS executives has become more apparent with the creation of a new title in the industry: Chief Information Officer or CIO. Recent surveys have indicated that a background in MIS is not necessarily a requirement and, in fact, may preclude an executive being considered for such a post. Why is this? One CEO stated in an interview that MIS executives were ''too technical, too limited in vision.'' This same CEO and others stated that ''an understanding of the business'' was far more important. Other remarks were made such as ''they understand information management but they do not understand information.''

The perception that MIS executives are ''too technical'' comes from the way such people have been used. More often than not the MIS leader was responsible for the *implementation* of systems and he may not have directly participated in other aspects such as selection of the system or the identification of the system's boundaries. In such cases, the MIS exec may have found it necessary to become technical to the degree that he could guide

such activities where poor choices were made. This has become more true as the shift from decisions about systems selection has been moved from MIS to the end user. In those cases the decisions are management driven, usually without consideration for the technical aspects. The technical aspects are still important. But if status meetings are held and the MIS exec reports on problems, the problems may be purely technical in nature and the chief executives or board may simply not understand them. Thus the MIS exec, through no fault of his own, is labeled "too technical."

The MIS executive may have a better understanding of information than he has been credited with. We have long known the value of "stored" information, but the fact that much of this information was or is stored in hierarchical databases that are rigidly structured and is often difficult to retrieve is lost on the information requestor. With the advent of relational databases, it is becoming easier to make this information available to the user. Relational systems are still not widespread in use, but they are a strategic tool.

Technological Limits

Information system planning from the tactical sense is still tied to the availability of technology. If you need the capability to process 400 transactions per second from a nationwide network, your options are limited. Until some improvements are made, these options will not include relational database systems. It is difficult to manipulate data in a system that is structure sensitive. It may be nearly impossible for the end user. Thus we may have been labeled "too limited in vision." It is not that we function with blinders on, but the technology limits what we can do.

A strategic information plan must be put together totally without regard to technology. Here is where the non-MIS executive may have an advantage. The non-MIS exec is probably not at all familiar with any information technology, so his approach is going to be data only. The MIS exec, either consciously or subconciously, may take a data-and-technology perspective even if he tries hard not to. This is a problem because a data-and-technology perspective invalidates any strategic plan. *Any information plan that is truly strategic is tied only to business goals and cannot be linked to technology.*

An example of this is a decision support system. You cannot design a strategic decision support system. You can only put in place some mechanism to make data available. This mechanism cannot in any way structure how someone might view the data because you cannot anticipate every view required. You cannot anticipate every business problem. If you put into place any decision support system with a structure, you have already limited its viability. It can only be regarded as a tactical tool.

Information Strategy

Strategic planning for information revolves around identifying what data is important and little else. The flow of information and any changes to that flow that could be caused by a change in business directions must be mapped. Large amounts of strategic data must be available to enable decision makers the capability of fast reaction to markets or changing trends. The implementation of the plan will of course be tied to technology — it must be. However, the information strategy will stay the same no matter the changes in technology. The technology will not affect the nature of the data.

In the beginning, when MIS was referred to as "data processing," we were relegated to installing bookkeeping and payroll systems. The limitations of the technology were blamed on the implementors. As information management evolved, our role evolved into planning and implementation. As a group I believe we are excellent planners. However with the advent of the CIO and some trend towards placing non-MIS people in those positions, we may find ourselves relegated once again to implementation.

If there is a way to curb this trend, it is probably through educating the people choosing the CIO. In any senior leadership position, qualities such as leadership, aggressiveness and vision should be top requirements. The senior MIS exec will have to promote himself in a situation where the creation of a CIO position is on the horizon. Becoming involved in activities totally unrelated to MIS may be the key to a successful move into CIO responsibilities. This becomes more important as other areas of a business besides MIS are almost always under the umbrella of the CIO position.

I hope that MIS executives can stem the tide in current CIO selections because that opens the door to many opportunities for other people involved in information systems. Perhaps we can be perceived as strategic weasons ourselves.

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.

Eliminating CICS Storage Constraints

By Larry J. Lawler

Eliminating CICS storage constraints and improving system response time is an ongoing battle. There are inherent bottlenecks in the design of CICS that can be formidable enemies. Attacking with bigger guns — a larger CPU, more DASD, more channels — does not necessarily remedy the problem. Although increasing these resources may temporarily offer some relief, they do not provide the permanent solution.

In the past, IBM has made significant advances in the stability and flexibility of CICS; however, relatively little progress has been made toward the elimination of storage constraints. Problems still continue to exist in this area. This article will address the affect they have on system performance and response time.

Q. What is the most significant storage constraint?

In most installations, the Dynamic Storage Area (DSA) is considered to be the most restrictive storage area affecting CICS performance. The DSA is the area that enables CICS programs to be quasi-re-entrant by providing a separate storage and work area for each transaction in process. Consequently, by having their own related storage areas in the DSA, multiple transactions are able to share the same copy of the executable program code. Transactions only use the storage areas while they are active and then release them when the task ends. On the average, a transaction usually requires approximately 32K of the DSA; however, this amount will vary depending on the specific processing requirements.

Many of the CICS system services also require use of the DSA to provide unique storage areas for the specific service they are providing. For example, the File Control Program (FCP) acquires a File Work Area (FWA) or File I/O Area (FIOA) for each record read from a CICS dataset.

Other uses of the DSA include: temporary storage main; dynamically created table entries such as the Processing Program Table (PPT), Program Control Table (PCT) and Terminal Control Table (TCT) when using the CICS on-line definition facility, (CEDA); terminal input/output areas for BTAM and VTAM terminals; basic mapping support mapsets; and application programs that were not defined as resident in the Processing Program Table (PPT) or Application Load Table (ALT).

Q. What determines the size of the DSA?

The size of the DSA is basically determined by the amount of storage that remains after CICS initialization. Following is an example of the procedure that can be used to determine the amount of CICS storage allocated to the DSA. First, subtract the storage required for the MVS operating system from the original 16M address space. A remaining region size of 8M is usually quite common. From the 8M, subtract the storage to be reserved for OSCOR. For this example, assume 1M was specified. This now leaves 7M. Next, subtract the size of the CICS system modules that average approximately 2M including programs and table entries. Now there are 5M left from which comes the storage used for VSAM buffers and control blocks and BDAM I/O areas for files opened at CICS initialization time.

The size varies greatly from system to system, but for now assume that there are 50 files in the system amounting to an accumulated storage size of 2M. Now only 3M remains for the DSA. However, to improve CICS response time, some programs are specified in the PPT or ALT to be permanently resident in memory. For this example, suppose that 50 COBOL programs of 40K each are to be permanently resident. This takes another 2M from the private region storage area. Now all that remains is a grand total of 1M of storage that must be shared as the Dynamic Storage Area (DSA) for all transaction processing.

Q. Will adding more real storage increase the size of the DSA?

FETCH DESTROYS CICS STRESS

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

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

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

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

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

FREE 30 DAY TRIAL

CIRCLE #84 on Reader Service Card ▲

Adding more real storage will do nothing to increase the size of the DSA. The 16M restriction is an imposed limitation of the operating system and CICS system design. Therefore, more real storage will not relieve these restrictions.

Q. How does the number of application programs affect the DSA?

The more application programs that are permanently resident in memory, the less the amount of storage that remains for use as the DSA. On the other hand, the fewer resident programs, the more programs there will be contending to use the limited storage that does exist in the DSA.

When there is not enough DSA storage available to satisfy a CICS GETMAIN request, CICS begins to compress the program subpool storage in the DSA. CICS program subpool compression is the basic process of proceeding sequentially through the list of application programs in the PPT and releasing pages of the DSA storage that currently contain unused programs.

Response time can be seriously degraded if storage compressions occur too frequently. However, some confusion exists as to what causes the poor response during these compressions. The main cause is the overhead that results from the flurry of I/O activity that occurs when the programs are reloaded back into the DSA when they are next used — not the actual process of releasing the program subpool storage itself. The storage release process can normally be completed in the time it takes for CICS to do just one I/O operation. However, the reload process may cause literally hundreds of synchronous READ requests, one at a time, from the CICS load library (DFHRPL).

Q. How does the CICS page size affect the DSA?

The page size as specified in the CICS System Initialization Table (SIT) can definitely affect system performance and storage utilization. Since all program sizes are rounded up to the nearest page boundary, using a page size of 2K can help reduce the amount of wasted program subpool storage. For example, with a 4K page size, an 8.5K program will allocate three contiguous 4K pages for a total of 12K. This leaves an additional 3.5K of program storage that will not be used. In contrast, a 2K page size will allocate five contiguous 2K pages for a total of 10K, leaving only 1.5K of unused storage. In this case, using the smaller 2K page size instead of the 4K page size would result in a savings of 2K in the DSA.

However, it may not be advantageous to use the 2K page size if your system is already experiencing noticeable paging activity. This is due to the fact that MVS pages in multiples of 4K rather than 2K. What CICS considers two pages may be one MVS page, or to make matters worse, two 2K pages in CICS could possibly span across two 4K pages in MVS.

The recommendation is to use a 2K page size if CICS paging is not an issue on your system. However, if your system is already experiencing a concernable amount of paging activity, then a 4K page size is recommended. The page size will also affect the ANTICPG=nn computation described later in this article.

Q. How does the size of the DSA affect paging?

The larger the DSA, the more the amount of real storage needed to support

the virtual pages. The problem is that one page-fault in the DSA may cause the entire CICS region to wait until the paging operation is complete. For this reason, CICS provides an Anticipatory Paging option in the PCT entries.

Using the Anticipatory Paging option (ANTICPG = nn) allows CICS to notify the operating system ahead of time as to which pages will be used, thereby eliminating page-defaults for transaction storage.

Since this does not apply to linked-to programs or the mapset storage areas of a CICS transaction, only part of the solution to the CICS paging problem is provided.

Q. How does converting to MVS/XA affect the size of the DSA?

When MVS/XA was initially introduced, there was very little effect, if any, made on the size of the DSA. However, CICS releases 1.6.1 and 1.7.0 made some changes to allow the CICS TRACE table, VSAM buffers (but not control blocks) and Temporary Storage MAIN to reside in the XA address space. The storage area that was previously used to contain this information is now free for other uses such as increasing the size of the DSA. The typical amount of storage relief provided by these changes is in the range of 1M to 2M. Unfortunately, in most cases this is still not enough to keep up with the ever increasing load of user transactions and growing demands placed on the CICS system.

IBM then announced the new PL/I and COBOL II compilers that support XA addressing — if you are using Command Level applications and are willing to convert all of your programs to the new compiler formats and syntax. For many installations with millions of lines of code and purchased applications or older macro level programs, the storage benefits offered by these compilers may not justify the expense or the conversion effort.

Even if an installation can afford the time and expense of converting to Command Level and COBOL II, this still only provides a partial solution. Basic Mapping Support (BMS) mapsets are still required to reside in the non-XA address space and, in many cases, a Command Level program may have a mapset that is as large as the program itself.

Furthermore, the COBOL II compiler uses operating system storage for COBOL Working-Storage areas instead of acquiring storage from the DSA. Even though this allows more storage to be available for the DSA, it also requires additional system overhead to acquire Working-Storage for each COBOL II transaction. Until now, this function was handled quite efficiently by CICS without the extra overhead. Now system performance is constantly being disrupted to process the SVC interrupt for Working-Storage GETMAIN/FREEMAIN requests.

Q. What can I do to reduce storage contraints in the DSA?

The main point to keep in mind when it comes to the DSA and CICS program loader is that this is a software-created problem. Increasing the CPU power or adding more memory is only a cosmetic solution at best. Tuning CICS in cases like this is also commonly referred to as "robbing Peter to pay Paul" because the resource used to improve one area must be "stolen" from another.

One alternative is to use the Multiple Region Operation (MRO) feature to create separate CICS regions such as one per application area. For example, payroll would be in one region, accounts payable in another and so on. This option results in more virtual storage being available because of fewer programs and less activity in any one region. However, the disadvantages that accompany MRO are the increased CPU overhead required to support multiple CICS regions and the complexity of maintaining separate on-line systems.

With each CICS region constantly scanning the dispatch queues for work to perform, it is not uncommon for an "idle" MRO region to utilize five percent of the CPU. If you multiply that five percent by the number of separate CICS regions you expect to keep active at any one time, you can see there is a definite potential for an excessive amount of wasted CPU power. Also, this does not take into account the fact that real storage consumption increases for each of the additional CICS regions.

For instance, one CICS region may have a 3M working set; a smaller region may have a minimum of 1M. Multiplied by the number of active CICS regions, the total working set size of real storage increases noticeably more than what would normally be required for just a single CICS region. Also, if Macro Level programs are involved, they cannot be shared between the separate MRO regions since they do not support function shipping.

Q. What are possible solutions to these problems?

For those who are presently experiencing DSA storage constraints, several different methodologies have been used to resolve the problem. For instance, one approach places the CICS programs in another address space and uses cross-memory services to retrieve them. However, because of its inability to handle page-faults efficiently and the increased demands placed on real storage, this approach is not a very practical solution. There is also the fact that there is still a storage limitation of approximately 8M when using the cross-memory mode, compared to two gigabytes available in the XA address space.

There are also other approaches that place the CICS programs on VSAM datasets to be reloaded each time they are referenced. The concept of providing multiple strings to the "simulated"

VSAM load library instead of the "single-threaded" CICS load library may appear interesting at first, but the CPU overhead of VSAM I/O is much too excessive to sustain efficient throughput.

The concept that I prefer is one that loads all CICS programs into the XA address space as they are first used; including Command Level, Macro Level, COBOL, RPG, Assembler, BMS mapsets and even 4GL programs. By loading the programs and mapsets into the XA address space, the size of the DSA is automatically increased proportionately by the size and number of resident programs placed there. (For MVS/SP systems, the same concept applies except that the programs are loaded into an "Extended Region" area by use of a VIO dataset.) This feature not only prevents the need to "split" CICS regions in order to obtain a larger DSA, but also it permits existing CICS systems to be combined for consolidated support and maintenance.

During CICS processing, the operating system is informed of the CICS program that is needed for processing the transaction and the anticipatory paging feature is used to page the program into the DSA. Then, at end-of-task the program storage in the DSA is released for use by other transactions.

Not only does this method significantly decrease the time required to load the CICS programs (since it is a memory-to-memory load and not a library-to-memory load), but also it provides a re-entrant-load process that allows multiple-load operations to be performed concurrently. It also ensures that the program pages are in real memory before the task begins to execute and thereby reduces page-defaults, eliminates subpool compressions and alleviates the DSA storage constraints.

It is not the intention of this article to imply that all CICS installations are experiencing serious DSA storage problems. If your system is not, then yours is among the fortunate ones. However, if you are about to place additional demands on your system in the future or your system already has this problem, then you may want to further research some of the previously suggested solutions.

ABOUT THE AUTHOR

Larry J. Lawler is vice president of product development for the software division of Quantum International Corporation, Birmingham, AL (800) 542-7760.

CCWTRACE Keeps An Eye On VM's I/O

By Ed Sterling

CCWTRACE is a very useful software tool available from IBM that provides highly detailed tracing of the real hardware devices attached to the mainframe. The trace information helps to diagnose difficult hardware and timing problems and gives both the customer and the vendor a permanent record of the interactions between the VM operating system and the hardware. Although VM provides Program Event Recording, virtual tracing and a CPTRAP command, CCWTRACE is in a class of its own. Unfortunately, because it is not a part of the regular VM/SP product, it is unknown to many new VM systems programmers. In this article, I hope to make you aware of the benefit of installing this utility on your VM system.

I have been a VM systems programmer for a long time and one of the major problems I have faced in dealing with errant hardware problems is the lack of a comprehensive and reliable data tracing tool for VM/SP. One of the most frustrating problem areas in VM is data communications, expecially the native ASCII and remote 3270 BISYNC support, handled directly by VM's Control Program (CP). Virtual machines that handle telecommunications lines, such as PVM (Pass-Through Virtual Machine) and RSCS, have very good internal trace facilities that allow the individual links to remote sites to be selectively monitored. VM/VTAM, of course, is a different "beast" and allows the use of VTAM and GTF tracing for the various SDLC lines, Physical Units and Logical Units. However, when hardware is directly controlled by CP, it *can* be quite difficult to diagnose and to isolate an input/output (I/O) problem.

In my pre-Simware days as a VM systems programmer, I developed software tools and CP modifications to collect and present information proving "my case" to hardware vendors. Now as a vendor myself, I have seen major flaws in the VM operating system bring down a customer's computer: the customer then blames me because the system failed only when using my software product! Yet ultimately, it was a "bug" in the VM Control Program that only came to life when my software product was run. The onus was on me to find the other guy's bug and prove that to the customer, quickly and under a lot of pressure! By having extensive event and data tracing information as well as some of the tools I will describe, I have been able to find those tough problems and assist my customers in solving them.

PER Helps VM Debugging

As far as tracing program flow for debugging, one of the finest productivity tools for programmers in VM is the Program Event Recording (PER) facility. PER itself is a hardware feature of IBM 370-class mainframes; but, it has to be activated and controlled by software so that each user or task can take advantage of it. As implemented in VM/SP, PER allows single-step instruction execution, tracing of certain instruction types over specific address ranges and also tracing of specific values of data stored in registers or specific memory location ranges. In a VM time-sharing environment, CP must control PER to ensure one user's use of it does not affect another userid. The PER facility has its origins at the University of Maine that for years kindly shared this complex CP modification with many other VM sites. Apparently IBM used PER (I assume IBM's own internal version) on IBM's development systems, yet it was not part of the general VM product until VM/SP Release 3.

CP's Internal Trace Facility

Although PER is great for tracing the logic flow in user programs, it cannot trace the real CP operating system and thus it is of little value when dealing with input/output problems. Looking into CP itself, one finds that all major events are recorded, however briefly, in the CP trace tables. A VM systems programmer who has read enough CP dumps gets to know the trace codes and formats quite well. Each storage "get" and "release," each scheduler call, each internal SVC and both start-I/O and I/O interrupts are all recorded in the trace table. There are more than two dozen different codes. This article addresses what the I/O trace entries can offer. Hex OB codes trace start-I/Os (SIOs) in which a list of channel command words (CCWs) is given to a device for various I/O operations. Hex 18 is a SIO fast-release, a variation of SIO that assumes the hardware devices are not likely to be already busy. Hex 05 is an I/O interrupt code that indicates the I/O request (OB or 18) is now complete and the hardware is signaling completion to CP. There are several other codes for test-I/O, halt/I/O, clear-I/O and so on. The trace entries offer limited information: the real device address, the condition code (important for SIO), the channel status word that gives some indication of why an I/O interrupt comes back with an error. But there is no data, nor any list of the CCWs to tell us what was being read or written to the device.

The dilemma in relying on CP trace codes is that either you have to take a CP dump (force an ABEND) or you have to work *very* fast to dump the CP trace area in which these entries are being recorded. Since they exist only in memory, they often have a lifetime of one second or less depending on how active your VM system is and how large your trace table area has been declared. Products like VM Systems Group's (Arlington, VA) V/SNAP allow a CP dump to be taken without bringing the system down which is a possible compromise. However, a CP dump represents a large and complex single snapshot of time. The real requirement is to have selected CP trace entries kept in a private area where they do not disappear. In VM/SP Release 2, IBM implemented a CPTRAP command that allows specific CP trace entries for certain userids and devices to be collected in a spool file. Yet even with CPTRAP, you still do not seem to have enough data to solve most hardware problems: there is no CCW list, no time-stamp and no data related to the I/O operation. For example, you do not know what the I/O operation was (such as POLL, SEEK, SENSE, ENABLE and so on and you have no idea what is being read or written.

A Close Look at CCWTRACE

Fortunately, there is an answer for this problem. It is an IBM "on-request" facility called CCWTRACE. Developed by IBM Canada, it consists of several modifications to the VM Control Program that collects detailed I/O information, as well as a CMS program that formats and prints the collected data for the systems programmer. You can obtain CCWTRACE from your IBM support center and it comes with the usual caveat of "use at your own risk." However, there is virtually no risk in installing or using it. CCWTRACE is not yet part of the official VM/SP product although it may eventually be included in Release 6 or 7. The good news for VM/XA/SP is that it has a DATATRACE command that does essentially the same thing as CCWTRACE.

CCWTRACE is a privileged CP command, activated by the systems programmer as required, that takes control of the CP tracing systems. It turns off normal tracing and uses the trace tables for its own data-collection area. Thus, one drawback is that if your system "abends," the normal trace data in the dump will not be there if CCWTRACE were running. The systems programmer can trace one or more devices along with a variable amount of data tracing. In some cases, your hunt for a bug may require maximum data tracing in the case where bad data is going to the device. At other times, you may be concerned more about the actual I/O operations in a case where there may be missing interrupts or exceptional status from the device.

Once tracing is completed, the data can be formatted and printed using the CCWTRPRT CMS module that locates and reads the real CP trace tables. The printout will show the device address, a millisecond-accurate timestamp, the channel command words and then a variable amount of data pertaining to the I/O operation. The CCWTRACE output can then be discussed by both the vendor and the customer to review the validity of the I/O operations, the content of the data

and timing aspects of the I/O operations. It becomes a definitive record of device operation and can prove a point for either side. In case you were wondering, CCWTRACE takes no measurable overhead during operation. My estimate is that it adds about 300-500 extra instructions per event which is really insignificant on a 30XX class processor. By comparison, simply pressing your ENTER key on a 3270 generates thousands of instructions in the VM operating system.

Normally, the CCWTRACE data is placed in the CP trace tables in a circular fashion wrapping over and over based on the number of devices being traced and their activity. For a problem that can be recreated easily, it is adequate to start CCWTRACE, produce the failure, stop CCWTRACE and then run the CCWTRPRT formatting program. When the need arises to trace a device all day long (that is once a day at random, a device drops off-line for no reason), CCWTRACE has an option that allows the trace information to be written to a tape drive dedicated to CCWTRACE. CCWTRPRT can then be run against the tape instead of the in-memory CP trace tables.

Experimental On-line Tracer

For a number of years, I have been working on a variation of CCWTRACE that allows me to "install" a CCWTRACE-like facility "on the fly." This capability is useful from a vendor's perspective because I may have to "parachute" into a tough customer situation and CCWTRACE may not be installed on that CPU. The other capability I have given my program is the ability to extract and display the trace data "real time." Hardware datascopes trace data on telecommunications lines and require some preparation to set up, unless you are lucky enough to have an extensive network lab.

My "software datascope" takes the CCWTRACE facility one step farther by reading the trace data, formatting it and displaying it on a 3270 terminal as quickly as possible. The display of data is similar to that of a hardware datascope with input and output data streams shown visibly distinct, mnemonics instead of hex codes and so on. The value of such a utility is that I can literally "see the problem" on my 3270 terminal the instant it happens, not minutes later as with CCWTRPRT. This utility has been valuable in examining how various devices work and helps in protocol analysis and reverse engineering.

In closing, having CCWTRACE on your system is good insurance for future problems that may arise. CCWTRACE can help you and a vendor quickly isolate I/O problems and it helps the confusion and finger-pointing that sometimes result from difficult system problems.

ABOUT THE AUTHOR

Ed Sterling is a principal founder and past president of Simware, a leading data communications software development company with a broad product line for VM/SP, MVS, VM/SNA, the IBM PC and the Apple Macintosh. He was previously a VM systems programming manager at one of Canada's largest VM installations. Simware is located at 20 Colonnade Road, Ottawa, Canada, (613) 727-1779. Sterling can be reached at Simware's Boston office, (617) 779-6058.

CACHE MAGIC™

With Cache Magic, improved VM system performance has never been easier. Enjoy a dramatic reduction in batch run times, in some cases by as much as 50%! Accelerate online response times, and complete more transactions in the same, or less time—all without the cost of additional hardware!

Cache Magic works its magic by storing your most active DASD files in an in-memory cache. Here they are accessed at processor speed, effectively eliminating I/O bottlenecks. And because Cache Magic automatically determines which files should be cached, performance gains are guaranteed, without investment of time and personnel.

Why not spend ten minutes with one of our representatives to find out what amazing things Cache Magic can do for you? All you have to lose are batch overruns, user complaints, and expensive hardware upgrades.

SDI.®

20 Years of Software Innovation
1 (800) 537-1969 (415) 572-1200
In Europe: +31 10 436 12 77
1700 South El Camino Real, San Mateo, CA 94402

CIRCLE #76 on Reader Service Card ▲

Advanced VM Diagnostic Techniques

By Gabe Goldberg

If an analysis of a "post mortem" dump of CP or other VM component is required, there are various aspects of reading a dump to understand. Understanding dumps' origins, contents and structure, being motivated to analyze dumps, learning about resources available in operating systems and identifying and acquiring dump reading skills have been addressed in the January/February 1988 and March/April 1988 issues of MAINFRAME JOURNAL. This article presents "post graduate" techniques for use when the standard diagnostic dump does not provide enough information for identifying and resolving a problem. It also outlines procedures for reporting problems to software vendors.

Sources of Additional Information

Static Information

Source Code

Source code is the form in which a program (application, utility or operating system module) is written by a programmer. It is the ultimate authority on what a program does, as opposed to what it is intended to do or documented as doing. It is also the preferred medium for program maintenance (rather than updating object code), because when used with standard VM maintenance tools and procedures it provides a reliable audit trail of changes.

Source code distribution and maintenance allows users to implement product enhancements and repairs as needed without being dependent on vendor priorities or abilities. When problem symptoms implicate a particular area of software, consult the source code (in listing form, of course, with current updates applied) to determine the logic and function being executed. Remember that occasionally comments in source code do not reflect actual program logic.

Load Map

IPLable components like CP, CMS and RSCS are built from many TEXT files (modules) during the system generation process. The VMFLOAD command builds an IPLable nucleus from modules listed in a "load list" (for example, "CPLOAD EXEC" for CP). The first module in a nucleus is an IPLable loader (DMKLD00E for CP) which reads the modules that follow, resolves cross-module references, produces a load map of the nucleus produced and (finally!) writes the nucleus on the DASD volume from which it will be IPLed when used.

The load map produced by DMKLD00E (or other component's loader) lists important information about each module in a nucleus (ordered by module load address):

Module name

- Filetype (indicates functional level assigned by the VMFASM command from the CNTRL file used for assembly so will not always be TEXT)
- Filemode and label of minidisk on which module was found during nucleus build process
- File timestamp (typically date and time of module assembly but may have been changed by copying file)
- Fileid, disk label, creation date and time of each update (PTF or fix) applied to module with associated AUX file comment
- Size (bytes of object code) of each module and the address at which it was loaded (and will execute)
- Address at which each external symbol or entry point was loaded
- For CP only, the boundary between resident (non-pageable) and pageable modules
- And the use of explicit or implicit Set Page Boundary (SPB) cards to force modules to begin on page (4K) boundaries.

The load map is an important tool for both traditional (paper, pencil, hex calculator) and automated (V/QUEST from VM Systems Group, programmable and command driven) debugging. It answers both dump-related questions (what module is at a particular address, where is a certain entry point) and system-related questions (which version of a module was used in the nucleus, what fixes have been applied to the module implicated by a dump). In its raw (print image) form, it is most effective mapping address-to-module (since it lists modules in order of load address) and weakest at mapping module-name-to-address (since that requires reading the entire load map to find the desired module). However, once the load map has been added to a dump with the VQUEST MAP command, both mappings are simple: the MODULE command identifies the module containing a specified address or the address of a specified module or entry point.

Manuals

Manuals are available from IBM and other sources relevant to dump-reading and problem diagnosis.

Dynamic Information

"Traps" and Diagnostic Code

When a problem is reported to IBM, the response is sometimes a "logic trap" or diagnostic code to insert in a suspected module to define new or more rigid bureaucratic rules. A trap typically causes an ABEND when a tested condition is detected in order to "snapshot" conditions as they are at the moment of failure, while diagnostic code may record extra information in a control block or display a new or augmented error message to identify the problem.

Enhancing and Freezing the Trace Table

As described earlier, some VM components automatically maintain a history of events called a "trace table." While usually of great value, a trace table only records events chosen by system developers. If a problem requires information not saved in the trace table, consider defining new trace table entry types to record needed information. It usually requires little code to place an entry in a trace table; macros or subroutines are often available to do this. New CP trace table entries have been used to record "sense" start I/O operations and associated sense information and stack/unstack CPEXBLOK operations. If a control block field overlay is causing a problem, the field can be saved repeatedly in the trace table; this will narrow the interval in which the overlay can be occurring.

After an event of interest has happened, a wraparound trace table may destroy needed information before it can be examined or saved. The CP trace table is controlled by bits at address X'400' with each trace table entry type controlled by a specific bit. When researching a problem that does not require all trace table entry types, the speed at which the table wraps (overlays previous data with new entries) can be reduced by only collecting needed data by turning off bits for unneeded entries. Less disruptive than forcing an ABEND dump to preserve the trace table when a trigger event occurs is stopping trace table logging so it can be examined later (perhaps with the CP option of VQUEST).

Trace Table and Supplemental Information

IBM CPTRAP and TRAPRED Commands

The privileged CP command CPTRAP collects trace table information, other CP data and user data in a spool file for analysis. CPTRAP spool files are class V, type DMP with fileid "CPTRAP FILE." Information can be restricted to specific trace entry types, I/O devices or VMBLOKs. A CPTRAP spool file can be specified to wrap, which limits both its size and the amount of data that will be available or not to wrap, which will produce spool files with a maximum size of 3480 4K records (pages). CPTRAP in non-wrap mode can produce large amounts of data very quickly; it should be monitored to prevent spool space from being exhausted. For example, the commands:

```
CPTRAP 6 VMBLOK 65AC0
CPTRAP START TO SYSTEST WRAP 200
```

collect only trace entries for FREE STORAGE (X'06') performed for the user with VMBLOK at X'65AC0' and directs a 200 record wraparound spool file to user SYSTEST.

The TRAPRED CMS command displays information from CPTRAP files which are non-standard format and cannot

be processed as normal spool files. For example, the commands:

```
TRAPRED 537
TYPENUM CODE 06
```

process spool file 537 displaying only code X'06' entries (which would be useful if the spool file contained a variety of entry types).

V/QUEST Commands

TRPDUMP and DMATRAP

The V/QUEST command TRPDUMP creates an IPCS/E minidisk dump file from a CPTRAP spool file or from a CCWTRACE tape. Extensive selectivity of data is provided; once the file is created, the V/QUEST TRACE command can be used to examine the data. For example, the commands:

```
CPTRAP START TO SYSTEST WRAP 200
CPTRAP ALLOWID SYSTEST
DMATRAP
#CP EXT
CPTRAP STOP
TRPDUMP
VQUEST TRP00001
```

collect (by default) all trace entry types with supplemental information from user SYSTEST, create file TRP00001 DUMP and process it with VQUEST. The V/QUEST TRACE and TRC commands provide better extraction and logging facilities than TRAPRED.

CCWTRACE

When I/O problems occur that cannot be diagnosed with standard tools, the IBM Support Center will provide a tool called CCWTRACE. This installs as a CP enhancement which requires a system generation. Once it is installed, it traces I/O activities and events in much more detail than the standard system; the additional information can help identify I/O problems in hardware or software.

V/SNAP and V/SAFE

V/SNAP and V/SAFE are VM Systems Group products which, respectively, take a CP dump without an ABEND and avoid the system outage normally associated with CP ABENDs.

Virtual Machine Debugging Tools

Several CP commands are available for debugging software executing in a virtual machine (which includes CMS, application software, RSCS, PVM and CP itself). In addition, many VM components provide self-debugging commands for gathering data. Frequently used virtual machine debugging commands are:

ADSTOP — Halts virtual machine execution just before an instruction at a specified address will be executed. Only one ADSTOP may be in effect at a time and ADSTOPs are cleared after being encountered.

DISPLAY — Displays contents of virtual machine components on terminal (storage, storage keys, general registers, floating-point registers, control registers, program status word, channel address word, channel status word).

DUMP — Prints contents of virtual machine components on terminal (storage, storage keys, general registers, floating-point registers, control registers, program status word, channel address word, channel status word).

PER — Monitors events during virtual machine execution. These include instruction execution, successful branches, alteration of a specific

general register and alteration of specified virtual storage areas. PER allows symbolic manipulation of event definitions.

STORE — Alters virtual machine registers and storage.

TRACE — Traces specified virtual machine activity (SVC interrupts, program interrupts, external interrupts, I/O interrupts, all instructions, privileged instructions, SIO instructions, branch instructions, CCW execution, CSW at I/O interrupt). TRACE is not as powerful or flexible as PER.

VMDUMP — Dumps virtual storage to a special format spool file for examination with a tool such as V/QUEST.

CMS debugging facilities and commands are:

EXECs — The three CMS EXEC processors (REXX, EXEC 2 and CMS EXEC) have integrated execution tracing facilities. Normally, an EXEC must execute the appropriate command to begin diagnostic tracing, but the CMS command

```
SET EXECTRAC ON
```

will trace the next EXEC 2 or REXX program invoked.

DEBUG — Enters a specialized debugging environment in which a number of subcommands are available. DEBUG is obsolete and vastly inferior to PER.

SVCTRACE — Produces a detailed trace of SVC execution including register contents before and after the SVC, name of called routine and location from which it was called and contents of SVC parameter list.

VM/Pass-Through debugging facilities and commands are:

AUDIT — Configuration file option to request logging PVM console data in specified CMS file.

DUMP — Configuration file option to specify whether PVM failure dump should be in VMDUMP or CP (print image) format. V/QUEST requires a dump to be in VMDUMP format.

QUERY — Displays internal PVM information including resource status and error counters.

SNAP — Snapshots internal PVM control blocks in CP DUMP format.

STATUS — Displays status of PVM system pools or tasks.

TRACE — Controls communication line tracing.

RSCS debugging facilities and commands are:

DUMP — Configuration file option to specify whether RSCS failure dump should be in VMDUMP or CP (print image) format. V/QUEST requires a dump to be in VMDUMP format.

QUERY — Displays link, file or RSCS status information.

TRACE — Controls communication line tracing.

Real Hardware

As a last resort, when debugging in a virtual machine (running CP second level or VM-under-VM) cannot solve a problem, processor control facilities can be used to trace execution or stop when a certain storage area is referenced or changed. This requires (in addition to a dedicated system) manually entering test conditions and recording addresses and other test results.

Reporting Problems

Problem analysis culminates in fixing a problem or reporting it. Problem reporting requires skill and finesse; simply telling the programmer next door or a vendor's customer support staff that "it doesn't work" will not usually result in a fix.

The more complete the symptom descriptions and documentation provided, the more likely the support person will be able to evaluate and fix the problem. Whenever possible, only copies of data should be transmitted — originals should be retained for local analysis and (just in case) for backup. A checklist for problem reporting includes:

- Symptom description (abend, loop, wait, incorrect output, message)
- Specific information (PSW for wait, sequence of addresses for loop)
- Supplemental information (CPEREP report for problem caused by I/O errors, VMMAP output for performance problem)
- Machine readable load map and dump of failing component
- Spooled console showing events prior to, during and after failure
- Information on local modifications
- Name of suspected module
- Real and/or virtual machine configuration
- Version, release, maintenance level of failing component
- History of product explaining whether it executed properly previously or has never executed
- Information on recent changes introduced
- Test sequence to reproduce problem
- Evaluation of severity of problem
- Customer number, name, authorization number and so on
- Contact person for additional information.

After reporting a problem, any interactions with the support organization should be logged with the date and names of people involved along with information received from the support organization, any problem number assigned, requests for additional data, data supplied and expected customer or vendor actions.

ABOUT THE AUTHOR

Gabe Goldberg is director of technology with VM Systems Group. He is responsible for product development and support and internal operations. Also, he is an experienced dump analyzer. VMSG develops and markets VM data center and end-user software.

Vital Signs *from page 49*

Unattended Operations

BlueLine is well aware of the interest its users have in automation issues and is, according to a spokesman, taking a careful look at how Vital Signs can address unattended operations. Along these lines, the company is currently beta testing Seekmiser, a tool that analyzes seek data from Vital Signs' database to actually affect the movements of CMS files on disk volumes.

In a VM system with multiple guest machines, the cause of poor CICS performance cannot be determined by analyzing only CICS performance data. Version 3.1 of Vital Signs includes a batch interface that permits CICS data to be imported from Landmark Systems' (Vienna, VA) The Monitor For CICS. This capability allows true correlations to be drawn between VM and CICS. Using the Vital Signs Report Writer and Plot Generator, CICS performance data can be reported in conjunction with VM performance data enabling true correlations to be drawn between CICS transaction repose time and overall CPU utilization, DASD contention and paging activity.

> ***A menu-based system that may be used by relatively non-technical people will be welcome.***

Prospects Bright

With the popularity of VM, especially in the coming era of the 9370, the prospects for Vital Signs seem good. As 9370s are installed by the thousands, the number of systems needing performance monitoring will radically increase. At the same time, the pool of trained personnel available to monitor and adjust VM performance will be strained. Thus, a menu-based system that may be used by relatively non-technical people will be welcome.

The menus used by Vital Signs are designed to allow managers the same access to VM internal performance data previously available only to sophisticated VM systems programmers. To further aid the non-technical manager, the system provides color graphic displays that make clear exactly what is going on.

Vital Signs is sold on a site license basis with a one-time site license of $8,000. Other permanent and renewable site licensing options are available as is a 30-day no-obligation trial. For more information, contact BlueLine Software, 1500 S. Lilac Drive, Suite 340, Minneapolis, MN 55416. Telephone: (800) 826-0311 or (612) 542-1072.

ABOUT THE AUTHOR

John Kador is a free-lance writer and a frequent contributor to MAINFRAME JOURNAL.

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

SNA Network Users Get Network Management Software

VM Software, Inc. has announced it will market a new product line of network management software for SNA networks.

The products, Net*Edit, Net*Monitor and Net*Control, will allow an IBM data center to simplify SNA network management by using object-oriented graphic definitions. This unique technology will allow comparatively inexpensive systems personnel to define, maintain and monitor VTAM host software freeing up present systems staff for more efficient use of their time. Net*Edit will be available first, during the fourth quarter of this year.

VM Software, Inc
1800 Alexander Bell Dr.
Reston, VA 22091
(703) 264-8000

For more information
CIRCLE #200 on Reader Service Card

New Electronic Mail Product from Westinghouse

NC-MAIL is a new stand-alone electronic mail product from Westinghouse Management Systems designed for the VSE, MVS and VM operating systems.

NC-MAIL utilizes a VSAM central administration file, which can be shared with Westinghouse's MULTSESS and NCI software, and features a panel-driven administration and full menu-driven system.

The product maintains lists for system, department and user distribution and can maintain lists by nickname, if desired.

Westinghouse Management Systems
P. O. Box 2728
Pittsburgh, PA 15230
(800) 348-3523

For more information
CIRCLE #201 on Reader Service Card

Technical Knowledge Checker Updated

Bookman Consulting just announced an update to TECKCHEK, an expert system composed of a comprehensive battery of tests that determines the extent and depth of an individual's programming proficiency.

TECHCHEK is not an aptitude test which indicates whether a person may become or has the ability to become an excellent programmer, but a proficiency test that measures what a job candidate (or employee) can do.

Tests are available now for IBM JCL, VS/370 Assembler Language, COBOL, COBOL II, SQL/DS, CICS, VSAM, Cullinet's IDMS, PCDOS, PC Assembler, etc.

Bookman Consulting, Inc.
Software Development
67 Wall St. - Suite 2411
New York, NY 10005
(212) 819-1955

For more information
CIRCLE #202 on Reader Service Card

Boole & Babbage Introduces DB2 MANAGER

A new performance management product for DB2 has been introduced by Boole & Babbage. Called DB2 MANAGER, it offers installations significant on-line facilities to monitor DB2 activity within the MVS environment.

DB2 MANAGER provides real-time analysis of DB2 performance and early warnings when user-determined thresholds are exceeded. It also provides displays of resource activity and utilization, enabling system administrators, performance analysts, database administrators and console operators to maintain constant surveillance of the DB2 Subsystem and MVS resources.

Boole & Babbage, Inc.
510 Oakmead Parkway
Sunnyvale, CA 94086
(408) 735-9550

For more information
CIRCLE #203 on Reader Service Card

New Application Communication System for VM/CMS

RD Labs has announced the availability of RD/COMM, a product which allows the creation of VM/CMS multi-user, multi-processing or service machine applications in high level languages such as COBOL, PL/1, C, Pascal and REXX.

RD/COMM provides high level language interfaces to IUCV (Inter User Communications Vehicle), eliminating the need to use CP or CMS Assembler macros.

The functions provided by RD/COMM enable the application to manage all facets of communications in a straight forward and consistent manner, regardless of whether there are interrupts from remote virtual machines or local utilities. No modifications to CP or CMS are required as RD/COMM will run as a relocatable module in the user's machine or as a DCSS.

RD Labs, Inc.
3825 Atherton Rd.
Rocklin, CA 95677
(916) 624-5755

For more information
CIRCLE #204 on Reader Service Card

IMS Database Downtime Reduced With IMAGE Copy Plus

A new solution to the universal problem of IMS/VS database downtime, caused by the need to create image copies off-line and the lengthy database recovery process, is now available from BMC Software. The new product is IMAGE Copy Plus.

IMAGE Copy Plus reduces the time required to create IMS/VS database image copies which are essential to the recovery of a database that has been logically or physically damaged. It functionally replaces the IMS/VS Image Copy Utility and enhances the IMS/VS Database Recovery Utilities with faster access and the capability for concurrently processing multiple databases.

BMC Software, Inc.
P.O. Box 2002
Sugar Land, TX 77487
(800) 841-2031

For more information
CIRCLE #205 on Reader Service Card

CICS Abend-AID for DB2 Out From Compuware

CICS Abend-AID for DB2 is a new application programming tool for anyone involved in abend analysis in a DB2 environment. It completely eliminates the manual, error-prone task of DB2 abend analysis in both the test and production environments.

CICS Abend-AID for DB2 tells exactly when, where and why a transaction abend occurred, as well as recommended solutions. On-line diagnostics are immediately available in a clear, easy-to-understand format. It is a powerful expert system tool which intercepts, analyzes and provides complete diagnostics for all DB2 DSNC abends and non-zero SQL Codes.

Compuware Corporation
31440 Northwestern Highway
Farmington Hills, MI 48018
(313) 737-7300

For more information
CIRCLE #206 on Reader Service Card

Multi-Image Manager Combines Best of SIS and SDM

Duquesne Systems has released Multi-Image Manager, a product that combines the performance of SIS (Single Image Software) with the functionality of SDM (Shared Device Management). The resulting product, Multi-Image Manager, provides total resource management capabilities for a multiple CPU or multiple image data processing environment. It is composed of several components that enable easy and cost-effective sharing of DASD, tape drives and console facilities.

The components of Multi-Image Manager operate from a single control file. Each component helps simplify, protect and control data integrity, tape drive allocation and console operations. Multi-Image Integrity ensures data integrity and manages data conflict across all systems; Multi-Image Allocation expands the tape drive allocation services of MVS and eliminates accidental tape overwrites; and Multi-Image Console consolidates multiple message streams to a single console.

Duquesne Systems, Inc
Two Allegheny Center
Pittsburgh, PA 15212
(412) 323-2600

For more information
CIRCLE #207 on Reader Service Card

Dynamic Buffer Allocation Product Boosts VSAM Performance

Initial test users of HYPER-BUF, a new dynamic buffer allocation software product from Goal Systems, are reporting consistant improvements in VSAM batch processing time between 30 and 80 percent. HYPER-BUF is the first product ever to automatically allocate buffer space for VSAM files under VSE and MVS as well as CICS.

HYPER-BUF intercepts VSAM OPENS and dynamically allocates I/O buffers based upon the access type (random, sequential, etc.), storage available at the

time of OPEN and user defined rules. The new product supports VSE 1.3.5 and above, all releases of MVS/SP and MVS/XA and CICS 1.5 and above.

Goal Systems International
7965 N. High Street
Columbus, OH 43235
(800) 848-4640

For more information
CIRCLE #208 on Reader Service Card

EMC Ships Solid State Disk Subsystems

Users of 43XX through 3090 systems may now want to consider a new device called the ORION from EMC Corporation. ORION is a solid state disk device for IBM channel compatible computer systems.

Access time for the ORION is 0.1 millisecond making it the fastest solid state disk system currently available for the IBM channel. It supports transfer rates up to 4.5M per second for maximum performance. Other transfer rates are 1, 1.5, 2 or 3M per second.

The ORION is compact (12″ by 29″ by 26″) and requires no special room preparation or air conditioning. It attaches to an IBM or PCM Block Multiplexer channel using standard bus and tag cables.

Upgrades for the ORION are available in 16M or 64M units. Maximum capacity for one unit is 544M, but daisy chaining units, up to 2.72G of capacity, can be configured using a single channel.

EMC Corporation
Hopkinton, MA 01748
(800) 222-3622

For more information
CIRCLE #209 on Reader Service Card

Change Man Available From Optima Software

Change Man, an automated library change and configuration management system, was recently announced by Optima Software. It automates the migration of software changes from test through production while ensuring the relationship between the source code and executable modules.

Change Man, which runs under TSO/ISPF, also handles JCL, COPY, database control files and documentation. It interfaces with the more popular dataset security systems and is said to be easy to implement because it is compatible with existing environments.

Optima Software, Inc.
1010 Hurley Way, #300
Sacramento, CA 95825
(916) 646-3800

For more information
CIRCLE #210 on Reader Service Card

Two New VTAM Products From SofTouch Systems

VTAM-WINDOWS is a VTAM session manager with fully interactive windowing of different VTAM applications. A user can work with TSO, CICS and IMS all on one screen eliminating time-consuming log-off and log-on between applications. Other features are: panning and scrolling of data in a window, VSAM sub-tasking maintains throughput, user command coding, etc.

VTAM-EXPRESS performs terminal datastream compression at the VTAM level. With it, a user no longer needs a separate compression product for each different VTAM application according to SofTouch. It also provides both outbound and inbound compression for every application in the network and separate operating regions are not required.

SofTouch Systems, Inc.
8269 S. Walker
Oklahoma City, OK 73139
(405) 632-4745

For more information
CIRCLE #211 on Reader Service Card

New Relational DBMS Tools for VSE/SP Users

Oracle Corporation, developer and marketer of the SQL-based ORACLE relational DBMS, has announced the availability of ORACLE and its associated development and end-user tools for IBM Mainframes under the DOS/VSE/SP operating system. CICS/VS is fully supported for high-performance on-line systems.

Until now, the only true relational DBMS available for users of DOS/VSE has been IBM's SQL/DS. ORACLE not only offers a fully compatible SQL alternative, but also is said to provide superior performance and a full range of development and end-user tools.

DOS/VSE users will be particularly pleased by SQL*QMX, ORACLE's compatible superset of IBM's Query Management Facility, according to an Oracle spokesman. Users can produce applications and offer end users services through Oracle's fourth-generation language (4GL) and decision support software (DSS) tools.

Oracle Corporation
20 Davis Drive
Belmont, CA 94002
(415) 598-8000

For more information
CIRCLE #212 on Reader Service Card

First VM-Based Capture/Playback Facility Announced

Unicorn Systems has announced VMCICS/Replay, a capture/playback facility, that allows the developer to save a sequence of transactions (screens) and user responses which can be replayed at any time. This facility is extremely useful for setting up long-running test scripts for the purpose of validating application programs.

VMCICS/Replay operates in two modes: capture mode which saves a sequence of test scripts in a file; and playback mode which processes a previously saved test script file.

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

For more information
CIRCLE #213 on Reader Service Card

VM Performance Tool From VM1

VM1 has announced QuickCopy, a new VM performance tool which reduces disk I/O and CPU overhead. It also reportedly increases DIRMAINT efficiency, helps eliminate disk fragmentation and enhances overall system performance in most VM areas.

QuickCopy does not occupy space in the user area, so it can be integrated into applications such as FOCUS and SAS to help boost performance.

VM1
5250 W. Century Blvd., #608
Los Angeles, CA 90045
(800) 321-4861

For more information
CIRCLE #214 on Reader Service Card

Multi-Functional Software Facility From Sequel

ProTerm, the professional's multi-functional software facility, has been announced by the Advanced Technology Division of Sequel Corporation.

ProTerm provides MVS systems professionals with the tools and facilities that extend their proficiency. In addition to its own unique functions, ProTerm provides all the features and facilities of several single-function software tools including: quality assurance testing, multiple session management, split screen operation, record and playback of live transactions, immediate review of prior-used terminal screens, transaction tracing and comparing, etc.

Sequel Corporation
Advanced Technology Div.
209 SW 89, Bldg. B
Oklahoma City, OK 73139
(405) 691-1498

For more information
CIRCLE #215 on Reader Service Card

New Integrated Journal Management & Recovery Program

FILESAVE/RCS, a completely integrated journal management and recovery program to reduce auxiliary disk storage requirements and recreate lost or damaged files, has been introduced by On-Line Software International. It is designed to assist systems programmers, data administrators and other operations personnel in developing a complete recovery system for managing both on-line and batch program journals, as well as performing forward or backward recovery of partially corrupted files.

FILESAVE/RCS provides users with a comprehensive Journal Management System for both CICS and batch program journals and a Recovery Control System. It derives its backward recovery capabilities by processing the "before images" of dataset changes.

On-Line Software International
Fort Lee Executive Park
Two Executive Drive
Fort Lee, NJ 07024
(800) 592-0009

For more information
CIRCLE #216 on Reader Service Card

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.

Disaster Recovery *from page 92*

Notification outlines the necessary notification procedure for the recovery teams. For example: damage assessment team and data control recovery team.

Assembly of recovery team contains instructions and procedures for the recovery team to meet after the disaster.

Assessment of damage/impact indicates the evaluation procedures to assess the impact/damage of the disaster on computer equipment, data information, personnel and services.

Recovery decision involves which plan will be executed. For example: on-site recovery or hot-site recovery.

Plan activation provides procedures for each type of outage.

Transition to hot site outlines how the data center should be generated at the backup site. Some key elements are: operating systems, network and component configurations.

Current Activities

The completion of a disaster recovery plan does not mean the end of all related activities. In fact, this is just the beginning since the plan needs to be constantly reviewed and updated and backup procedures and plans need to be changed according to changes made to the recovery plan. Also, test plans should be developed and exercised to ensure that the plans are functional and walkthroughs (dry runs) must be scheduled and carried out to ensure good attention to detail.

Conclusion

Much has been written about the technical side of disaster recovery. Contingency planners know that the technical details are easy to handle, that the most difficult yet most rewarding part of the work is the homework leading up to approval of the final plan and the efforts required to get the plan endorsed by senior management.

Tying a project of this magnitude together will be complex, dynamic and expensive. However, successful implementation of a disaster recovery plan can make a significant contribution to the uninterrupted data processing function of your company.

ABOUT THE AUTHOR

Kern Chang is a systems engineer for National Advanced Systems, Dallas, TX.

LEAK-SENSING DEVICES

Water incursion in a data center can be devastating. According to the National Insurance Institute, 35 percent of data center insurance claims result from damage caused by water. In fact, three times more water-related claims are filed than fire related. The following have all been traced as the cause of data center problems:

- Condensation on equipment caused by high humidity
- Plumbing leaks
- Sweating cold return water pipes
- Clogged drains causing overflows
- Thawing of frozen pipes in building infrastructure
- Chilled water unit leakage
- Rain entering data center from a breached structure
- Ruptured storage tank
- Leaks from sprinkler heads
- Accidental discharge of sprinkler system
- Overflowed toilet
- Janitorial maintenance accidents
- Moisture entering through cable cavities
- Storm-caused flooding
- Cracks in substructure and low frost line
- Flooding of data center by fire department.

There are three basic forms of sensing devices available: cable, point and tape. With a cable device, the coverage is continuous. The cable is laid around the perimeter of the subfloor and in a grid pattern under risk areas.

A point device gives spot coverage. One point is positioned every 500 square feet. It can be implemented singularly or as a network wired into a remote telephone dialer. Third is a tape device that gives zone coverage. It is placed in a box pattern under risk areas.

With each of these devices, ambient noise in the data center could make the alarm difficult to hear. Using a remote annunciator eliminates this situation.

Another important application for remote annunciators is the unattended data center. Do not get caught without backup power for your sensing devices. Several data centers experiencing water damage would have been able to lessen the damage if there had been an early warning. Unfortunately when the power went off, so did the sensors.

Leak-sensing devices should be placed near air conditioners, chilled water units, dehumidifiers, humidifiers and heat rejection equipment.

Properly locating your data center away from basements and top floors is at the foundation of avoiding water damage. However, emergency pumps and waterproof covers should be inventoried in your disaster recovery kit just the same. Leak-sensing devices will provide you with the precious minutes you need to save your data center.

By Tari Schreider, president. Contingency Planning Research, Greenwood Landing, N.Y.

THE IBM MID-SIZE MAINFRAME DISASTER RECOVERY SOLUTION

On-Site 3890 Capability

Cary, North Carolina

Tewksbury, Massachusetts

Niles, Ohio

Relocatable Shell Facilities

HOTSITE® offers:

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

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

HOTSITE®
A DIVISION OF COMPUSOURCE

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

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

Name _____

Title _____

Company _____ *Phone* _____

Address _____

City _____ *State* _____ *Zip* _____

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

MJ 9-88

CIRCLE #40 on Reader Service Card ▲

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

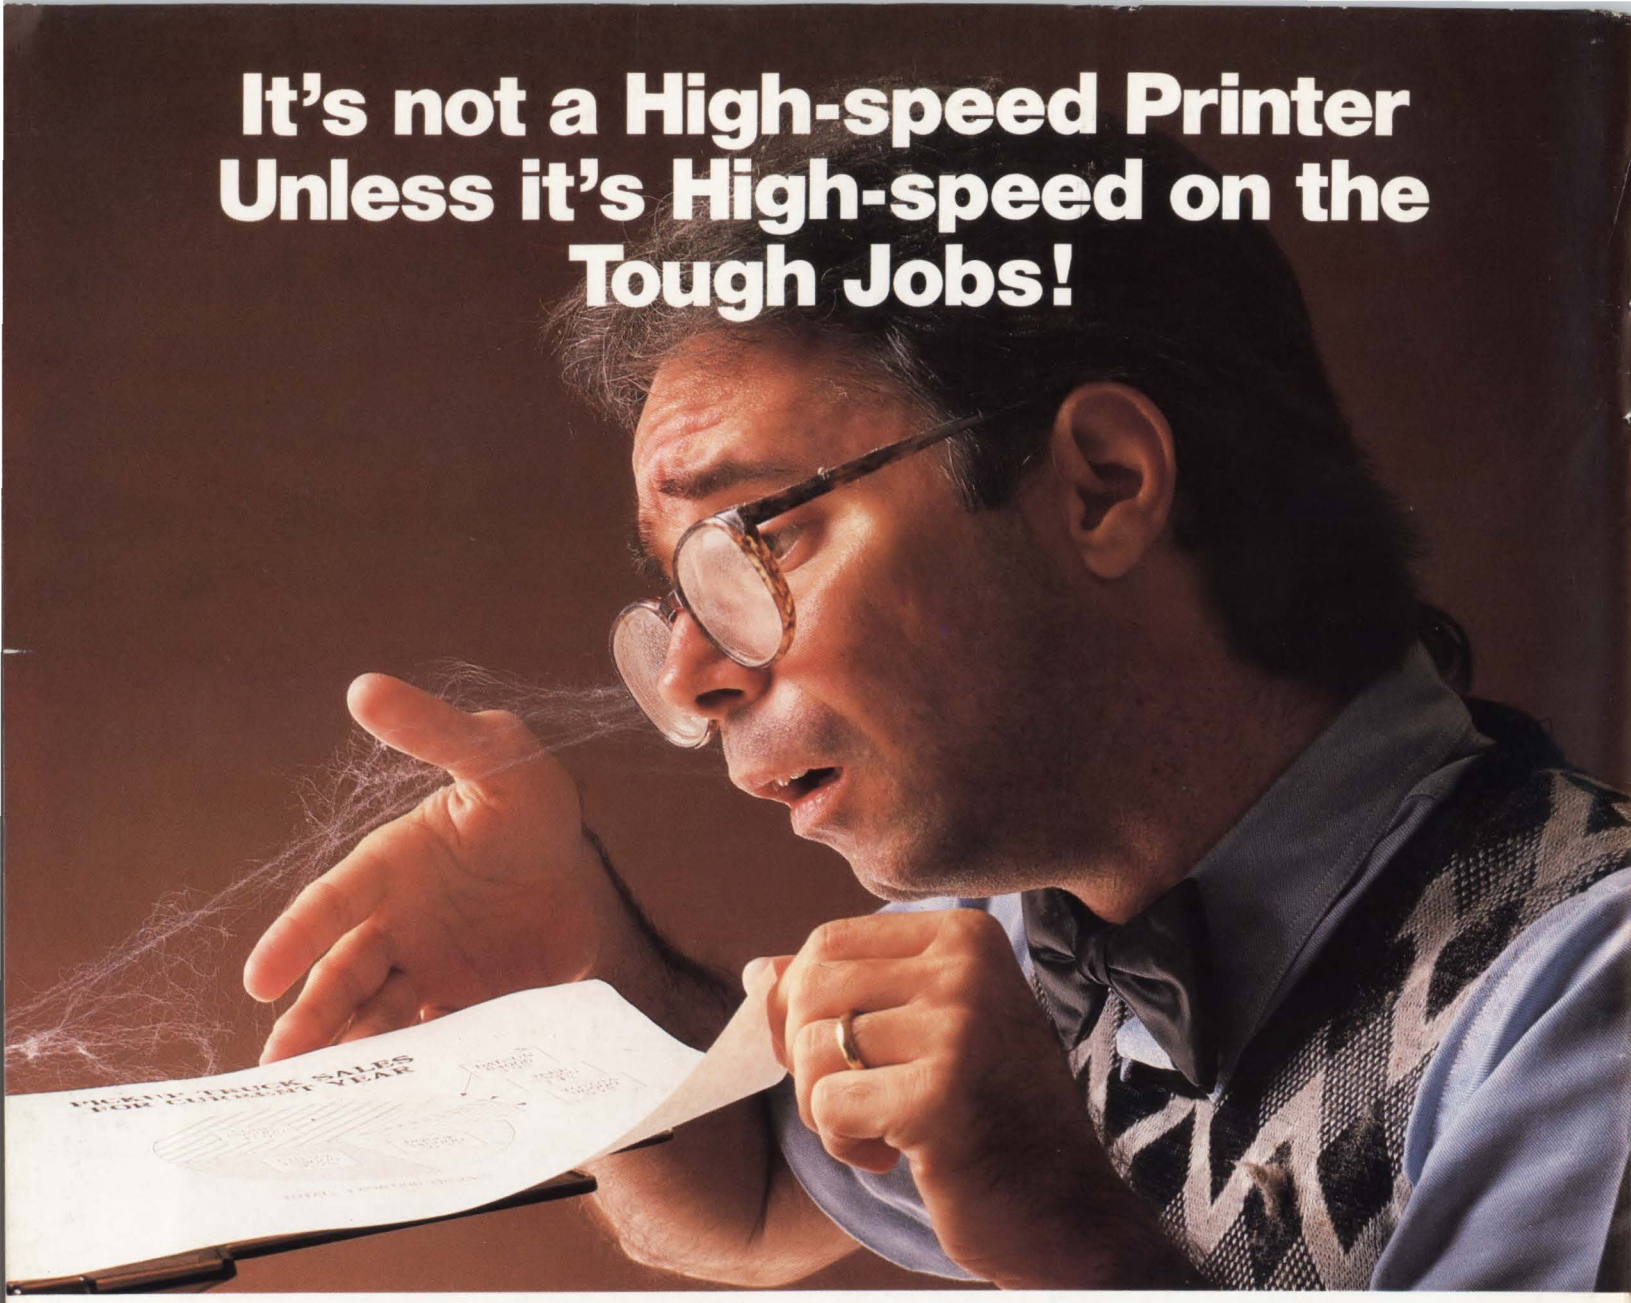

Introducing the NBS 3840 Laser Printer for a TRUE 40 pages per minute!

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.

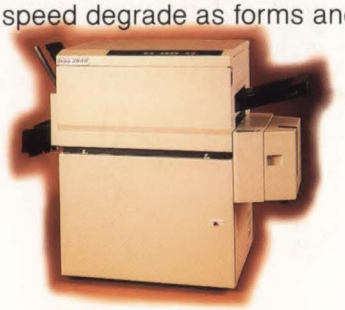

With the addition of NBS Southern's PAGEWARE™ Forms Description Language (FDL), you can easily merge text with complex documents, logos, signatures and other special graphics.

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.

CIRCLE #102 on Reader Service Card ▲