

Managing VM/CMS systems

Dear Editor:

I found the article "Managing VM/CMS systems for user effectiveness" (Volume 18, Number 1) by W. J. Doherty and R. P. Kelisky to be quite interesting. There are, however, several points in the article that I believe need to be corrected or clarified for the reader who is not familiar with the VM/370 systems at the IBM Thomas J. Watson Research Center.

Forum

Though it is not mentioned by name, many of the data base references in the article are to a facility called MIGRATE. The distinction within MIGRATE between migration (i.e., on-line storage management) and backup (i.e., copying user data for availability reasons) should be made clear, as should the differences between the TSS and VM/370 systems on the subject of one-level store. The authors refer to design considerations and early implementation features of MIGRATE, but fail to note that in practice several of these items (e.g., the dismounting of off-line volumes and the conditioning of the backup on the THI) have been found not to work as planned. Unlike the MVS Hierarchical Storage Manager (HSM), MIGRATE does not handle files but, rather, minidisks. There is a separate file archiving facility at the Research Center called ARCHIVE which aids users who wish to manage their own mini-disk space. It should also be noted that the AUTOLOGON facility is primarily an aid for the system operator at restart time to reduce the number of commands he has to enter. It does not guarantee the restart of MIGRATE or of any other subsystem which fails while VM/370 continues to run.

A major omission in the article concerns the use of the MSS (IBM 3850 Mass Storage System) with VM/370. The recently distributed Release 6 of the VM/370 Control Program includes support for minidisks residing on Mass Storage Volumes in the MSS. At the Research Center, we have been using the MSS with VM/370 for over three years. The reduction in tape and disk mounting attributed by the authors to MIGRATE is mostly due to the MSS installation. Backups of MIGRATE-controlled minidisks are done automatically through VM/370 to the MSS overnight. Backups of MSS-controlled

Copyright 1979 by International Business Machines Corporation. Copying is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract may be used without further permission in computer-based and other information-service systems. Permission to *republish* other excerpts should be obtained from the Editor.

minidisks are done entirely within the MSS on weekends using no central system resources.

David L. Reich
IBM Thomas J. Watson Research Center
Yorktown Heights, New York 10598

Authors' Response

Since the letter corrects our article as to several aspects of the VM/370 systems at the IBM Thomas J. Watson Research Center, we take this opportunity to discuss these matters a little more fully here. There is ambiguity in our use of the term "data migration." With respect to VM/370, the MIGRATE system provides the following two distinct facilities: (1) direct access storage device (DASD) management, which attempts to control space pre-allocated to user minidisks by moving user minidisks, not files, on line on demand, and (2) minidisk backup, which copies changed minidisks into the MSS (IBM 3850 Mass Storage System). The MIGRATE system, although prompted by our experience with TSS, does not give the user the appearance of an unbounded one-level store as does TSS. MIGRATE enables the system administrator to allocate minidisks for users as though there were more on-line capacity than is actually available. We have observed, however, a performance penalty that increases with the degree of overcommitment. Although MIGRATE was planned to reduce the need for on-line disk drives by moving less frequently used minidisks to demountable disks, in practice off-line minidisks are not demounted.

The criterion for migration of a minidisk to off-line volumes from on-line volumes is not based on a specified number of days since the minidisk was last used, but on a threshold of logical emptiness on the on-line disks. That is, in anticipation of new requests for migrated minidisks, a certain amount of on-line space must be available. Less frequently used on-line minidisks are marked as migratable, copied to off-line volumes in advance by nightly "look ahead," and then replaced on line when space is required by other minidisks.

Although MIGRATE once conditioned its backup schedule on the TH1, this procedure was changed when backup failed to finish. Backup is now begun about 6:00 p.m., and normally completes at 5:00 a.m. (It is terminated at 8:00 a.m. whether complete or not.) Migration and restoration proceed without regard to system load.

As the correspondent points out, the VM file migration facility, called ARCHIVE, is logically distinct from MIGRATE. The user is responsible for managing the space within his minidisk, but ARCHIVE enables the user to move files to and from that minidisk by means of simple commands. Minidisk backup provided by MIGRATE is also logically distinct from the minidisk migration service. The reference to backup to tape is historically correct. Tape backup, as described in the article, was done on one of our VM/370 systems until March 1979. On both VM/370 systems minidisk backup is now done automatically to the MSS. Recovery of a user's backed-up migratable minidisk now takes four to six minutes as compared with hours previously.

It should be noted also that availability of the MSS has encouraged users to move large blocks of data to private minidisks on virtual MSS volumes. As a result, the user still has access to his data even though tape and disk mounting are unavailable when the systems run unattended. MVS/TSO/JES3 and VM/370 no longer have separate operator teams although that was our practice until late in May 1978.

The description of the AUTOLOGON facility may have implied a guarantee that facilities such as MIGRATE are always available after a system failure. MIGRATE is not restarted by AUTOLOGON if CMS itself is down or if there are bugs in MIGRATE.

The article attributed our reduction in tape and disk mounts to the MIGRATE facility, and this is not entirely correct. Rather, reduction in volume mounting has been due to the introduction of the MSS for large user minidisks on Mass Storage Volumes, backups of MSS minidisks in the MSS, and creation of virtual MVS and OS/VS1 systems on Mass Storage Volumes. MIGRATE enabled some users to release backup tapes, and may have decreased prime shift tape mounting. But it is primarily the introduction of the MSS into our operational environment that has been the major cause for reduction in volume mounting being done in a period of continued marked growth (1976-1979).

Another factor that contributes to the reduction in volume mounting is that minidisks can be shared between our two VM/370 systems either by means of Mass Storage Volumes or permanently mounted volumes (but not by MIGRATE). As a result, duplication of on-line data has been reduced, as well as associated tape and disk mounting.

The major role of the MSS at the Research Center was not addressed in this article, and the reader is referred to two publications: (1) *Using MSS with VM/370 at the T. J. Watson Research Center* by David L. Reich, Research Report RC-6541, IBM Thomas J. Watson Research Center, Yorktown Heights, New York

10598, and (2) *A Manufacturing Company's Experience Using the IBM 3850 Mass Storage System to Store Interactive User's Data Sets*, IBM Application Manual, Form No. GE20-0627-0 (available through IBM branch offices), which describe in detail the way in which the MSS functions with VM/370 at the Research Center.

The authors acknowledge the contribution of D. L. Reich, who developed the facilities that support the use of the MSS for VM/370 minidisks at the Research Center, and regret their inadvertent omission of that acknowledgment in the original article.

Reprint Order No. G321-5112.