

# Staying on Track: DFSMS Version 1.4

BY STEVE PRYOR

The one constant in the modern data processing environment is change. As data centers continue to grow into the multi-terabyte range, the storage administrators and systems programmers who deal with this growth are faced with the daunting task of simply attempting to keep up with the latest developments. In a large data center there may be dozens or even hundreds of different products and operating system components at varying levels of maintenance. Ensuring that all of these components and subsystems work together properly can be very difficult in an environment of such complexity.

In an attempt to reduce this complexity and its burden on the managers of the system, for the past year or so IBM has been moving toward a scheme where not only the OS/390 operating system itself but all of its associated components and subsystems (e.g., DFSMSdfp, DFSMSdss, RACF, CICS, DB2, etc.) are kept "in sync" and distributed as a single unit. To illustrate, new operating system releases can be viewed as a "train" that leaves the "station" about every six months, with the various subsystems loaded onto the train's freight cars. Of course, installations are not necessarily expected to install an entirely new release of all of their systems every six months. However, since the contents of each "train" are to be tested together extensively, they are more easily installed, less complex, less subject to breakdown and can run more often. Decreasing the level of difficulty and expertise required to install and maintain the operating system is a key requirement for IBM to allow OS/390 act as an enterprise-wide server operating system. While not all subsystems are yet included in this program, the direction of more frequent, more integrated OS/390 releases is clear.

For the storage administrator, the most important "car" on the OS/390 train is DFSMS, and the most recent "departure" was on September 26, 1997, with the inclusion of DFSMS version 1.4 in OS/390 release 2.4.

Many installations have either already installed this release or are planning to install it, so this month I'll examine some of its most important features and how they affect the storage administrator's task.

**Backup and Restore:** Ensuring speedy backup and reliable restore is probably the most basic requirement that storage administrators face. A number of changes were included in DFSMS version 1.4 for DFSMSShsm and DFSMSdss that affect backup and restore operations. Among these are:

- ◆ *Duplex Tape* – A new SETSYS DUPLEX command finally provides DFSMSShsm with the ability to create two tape volumes at the time of backup or migration rather than relying on TAPECOPY to create additional tapes after backup or migration processing is complete. The duplicate tapes use the character string 'COPY' as the second qualifier in the dataset name.
- ◆ *ABARS Improvements* – DFSMS version 1.4 adds several enhancements to Aggregate Backup and Recovery Support, IBM's methodology for performing application-oriented backup. These include increasing the number of possible ABARS tasks from 16 to 64, reducing the number of backup tapes required by stacking all aggregate output files on a single tape, and improving performance by allowing the OPTIMIZE keyword to be passed to DFSMSdss.

- ◆ *Record Level Sharing for DFSMSShsm Control Datasets* – The DFSMSShsm control datasets (BCDS, MCDS, OCDS) can now be shared via VSAM record-level sharing, which takes advantage of coupling facility hardware to synchronize cross-system access. For large installations with many DFSMSShsm hosts, this may improve performance.

DFSMS 1.4 introduces a number of other changes to backup and restore, such as the extension of the DFSMSdss space-release function to logical processing and extended-format VSAM datasets, the new SELECT-MULTI keyword to control the selection of multi-volume datasets, and more.

**VSAM Improvements:** VSAM clusters are the most common type of dataset in many installations and thus are often of particular importance to the storage administrator. A number of changes have been made in this area, including:

- ◆ *Last-Reference Date Update* – Dataset migration via DFSMSShsm, FDRABR or other tools relies upon the last-reference date in the format-1 DSCB to determine if a dataset is "old" enough to migrate. Prior to DFSMS 1.4, this date was set only at the time a dataset was opened, which led to datasets that remained open for long periods to appear to be unused when they really were not. DFSMS 1.4 eliminates this problem by setting the last-reference date when the dataset is closed as well as when it is opened.
- ◆ *Extended-Format Enhancements* – Several improvements have been made to processing for extended format (striped) VSAM datasets, which of course can

exist only on SMS-managed volumes. These include new JCL parameters and data class attributes (RECORD ACCESS BIAS=SYSTEM) that allow the system to determine the buffering for VSAM datasets. Very large datasets, such as those greater than 4GB in size, which require extended-format, can now use the coupling facility for record-level sharing. The number of possible LSR pools has also been increased in DFSMS 1.4 from 16 to 256.

**Other Changes:** There are several other changes in DFSMS version 1.4 that may be of interest to the storage administrator. These include the following:

◆ *Catalog Search Interface* – This facility, which was described in more detail in last month’s column, allows a user to write a program to obtain information from the catalog for selected datasets using a generic mask such as ‘SYS2.AB?D.\*’. This is a great improvement in both performance and simplicity over other methods of accessing the catalog. A sample

REXX exec that invokes the Catalog Search Interface is included with DFSMS 1.4.

◆ *Space Error Retry* – DFSMS version 1.4 includes new parameters in the data class that will allow some types of out-of space conditions for SMS-managed datasets to be prevented by retrying the allocation with a smaller value or more extents. While only limited types of recovery are performed, it is a step toward the more complete facilities provided by vendor products that prevent out-of-space errors.

◆ *SMS Volume Selection* – DFSMS version 1.4 introduces changes in the way that SMS selects volumes at allocation if the primary list is empty and a volume must be selected from the secondary list. This support, which was previously available in APAR OW23333, randomly selects a volume from the secondary list instead of always selecting the volume with the most space, which may improve the placement of database files.

This column has only touched on some of the more significant changes introduced with DFSMS version 1.4. There are many others that may be important to storage administrators, including batch ACS routine testing via NaviQuest, improvements to dynamic cache management, DCOLLECT, and the DFSMSrmm tape management system.

The next “train,” OS/390 version 2.5, is already on its way out of the station and version 2.6 is being loaded. For the storage administrator whose responsibilities include getting the most out of enterprise system software (and that’s all of us), an understanding of what capabilities are available in DFSMS and how they can be exploited will help keep the corporate train moving ahead. 

---

**NaSPA member Steve Pryor has more than 15 years of experience in storage management, disaster recovery, software development, and technical support. Steve can be contacted via the Internet at [pryor@atlanta.com](mailto:pryor@atlanta.com).**

©1998 Technical Enterprises, Inc. For reprints of this document contact [sales@naspa.net](mailto:sales@naspa.net).