SMS Constructs Explained

BY STEVE PRYOR

If you were lucky enough to attend the recent joint SHARE/GUIDE meeting in Washington, D.C., you probably couldn't help but notice the celebrations that proclaim this year the 10th anniversary of IBM's introduction of system-managed storage. In the years since its introduction, system-managed storage has become the most important way for installations to provide effective control over the use of the vast storage resources required to support the corporate mission.

The heart of system-managed storage is the DFSMS architecture with its Automatic Class Selection (ACS) routines that supply policy rules for the management of data and the four basic SMS constructs (data class, storage class, management class, and storage group) that allow datasets with similar needs to be managed in the same way. The SMS constructs are critical to correct creation, placement and management of data, yet often only a handful of storage administrators really have a thorough understanding of their function and relationship to one another.

This month, I'll look at the purpose of each type of construct and briefly examine its structure; my goal is to provide a better understanding of how the constructs affect the use of data.

The SMS constructs provide a means of classifying data according to different types of criteria such as physical characteristics (data class), performance requirements (storage class), and backup availability (management class). The ACS routines provide a set of rules for determining what datasets belong to each classification and assigning datasets with those characteristics to DASD volumes (storage groups).

Almost every characteristic that could be associated with a dataset can be found in the SMS constructs, with each construct having anywhere from 15 to 40 or more individual attributes. The key to dealing with this complexity is to understand the specific function of each construct.

Over the past 10 years, the SMS constructs have grown in complexity as the amount of data under system management has increased and as new hardware capabilities have appeared. An understanding of the characteristics of these structures is important not only for the storage administrator, but for anyone who uses or manages company data.

THE SMS CONSTRUCTS

The Data Class construct is used at initial allocation of a dataset and supplies information about the type of dataset that is to be created, such as amount of primary and secondary space required, LRECL, DSORG, etc. VSAM dataset characteristics such as control interval size and sharing options can also be specified. Unlike storage class and management class, the system does not usually retain the name of the data class associated with a dataset since it is not needed again once the dataset has been created. Also, unlike storage class and management class, data classes can be assigned to non-SMS datasets.

Because the data class attributes RECFM, EXPDT, volume count and the like can be overridden in the JCL, the data class construct is not widely used for this purpose. Instead, its most common use is for SMS-only attributes such as DSNTYPE. This attribute allows the creation of extended-format datasets that support striping or PDSE (non-VSAM) or extended addressability (VSAM datasets more than 4GB in size). The COMPACTION attribute is also used to allow data compression.

The Storage Class construct is probably the most important construct from the DFSMS point of view, since, very simply, datasets that are assigned a storage class by the ACS routines become SMS-managed, while those that are not assigned a storage class do not, and must be managed in some other way. For SMS-managed datasets, the storage class provides a means of specifying the performance objectives for a dataset, primarily via the type of device on which the dataset will be created. Attributes for the desired response time in milliseconds (MSR) and the type of access the dataset will have (read/write and direct/sequential) allow the system to choose the type of device that is most likely to meet these requirements. SUSTAINED DATA RATE allows the system to create a multi-stripe dataset (if the Data Class specifies DSNTYPE=EXT). The AVAILABILITY attribute refers to the reliability of devices on which the dataset may be allocated — dual copy, RAID, or ordinary DASD, while the ACCESSIBILITY attributes control the selection of devices that allow continuous accessibility to data during backups, i.e., those behind Concurrent-Copy-capable controllers or on SnapShot-capable RVAs.

The Guaranteed Space attribute also belongs to the Storage Class construct and not, as is sometimes thought, to the Storage Group. Guaranteed Space, which has no

effect once allocation is complete, allows datasets that are assigned storage classes containing this attribute to be placed on specific SMS-managed volumes. This facility is most commonly used by database administrators who want more detailed control over the placement of DB2 datasets.

The construct that is probably of most interest to storage administrators is Management Class. It is by far the most crowded and complicated of the constructs, with five separate ISMF panels dedicated to its definition. Unlike the data class and storage class constructs, Management Class attributes come into play after datasets have been allocated. The major function of Management Class is to control the movement of data through the storage hierarchy via migration and backup, and to specify how long data is to be retained at various levels of the hierarchy.

An examination of the many different Management Class attributes shows that they are actually divided into several groups according to their purpose. The most commonly used attributes are those devoted to space availability management for primary DASD — that is, migration and associated functions. The PRIMARY DAYS NON-USAGE attribute indicates how long data may remain unused on DASD before it migrates to another level. LEVEL 1 DATE/DAYS indicates how long the migrated data may remain unreferenced before it is again migrated lower down the hierarchy. Other attributes are used to indicate whether older GDG generations are to be migrated or scratched, and whether allocated but unused DASD space may be released from overallocated datasets. The Expiration Attributes (EXPIRE AFTER…) indicate whether a dataset may be scratched if it is unused.

The Backup Attributes group provides criteria for how often backups (as opposed to migrated copies) of datasets are to be taken and how long they are to be retained, as well as whether Concurrent Copy should be used during backup.

Less commonly used are the Object Class Transition attributes that are associated with DB2 OAM objects, and Aggregate Backup Attributes, which are used for the control of ABARS (application-oriented) backup execution and the retention of ABARS backup tapes and logs.

Complicating the use of management class somewhat is the fact that the major DASD management systems, DFSMShsm and FDRABR, as well as others, interpret the characteristics of the management class attributes slightly differently. For example, setting ADMIN OR USER COMMAND BACKUP = NONE indicates to DFSMShsm that HBACKDS commands cannot be used to back up these datasets, while it indicates to FDRABR that these datasets are not to be included in TYPE=ABR backups. This difference in processing also applies to the Storage Group construct.

Storage Group is different from the other three constructs in that it does not apply to datasets, but rather to volumes. Several different types of storage groups may be defined (VIO, TAPE, OBJECT, and so on), but the most commonly used is, of course, POOL, i.e., DASD volumes. One of the more important attributes associated with the storage group is the allocation/migration threshold, a single number that applies to all of the volumes assigned to the storage group and which serves two purposes. The first is to direct new datasets away from volumes that are becoming full, thus spreading the allocation load across the volumes in the storage group. The second is to inform migration managers such as DFSMShsm and FDRABR when enough space has been freed up so that further migration on the volume can be bypassed.

Both storage groups and the volumes associated with them may have a status with respect to SMS (ENABLEd, NOTCONNected, QUIESCEd) that indicates whether they are available for use. While the ACS routines may assign multiple storage groups to a dataset during creation (SET &STORGRP='SGA', 'SGB', 'SGC', etc.) this simply provides a list of volumes for DFSMS to choose from when allocating the dataset. DFSMS and SRM, the System Resource Manager, order the list of eligible volumes according to algorithms that take into account the attributes of the storage class assigned to the dataset, device performance characteristics, free space on the volumes, and other factors. Once allocated, datasets reside on volumes within a single storage group and cannot extend outside of it. If additional space is needed within a storage group, volumes must be added to it via ISMF.

Besides the four primary constructs, there exists an additional construct, the Aggregate Group, which defines the characteristics of application-oriented backup units. Aggregate Group definitions include the names of "selection" datasets that provide a list of files belonging to an application, the number of backup copies to make, management class for the aggregate group, and so on. Aggregate groups are backed up and recovered by the DFSMShsm ABARS (Aggregate Backup and Recovery) function.

Over the past 10 years, the SMS constructs have grown in complexity as the amount of data under system management has increased and as new hardware capabilities have appeared. They have also grown in importance in allowing the massive growth of storage in recent years to be effectively controlled. An understanding of the characteristics of these structures is important not only for the storage administrator, but for anyone who uses or manages company data. ts

---

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.

