SMS Volume Selection

SHARE in Anaheim
Session 3020
February, 2005
Ruth Ferziger
ruthf@us.ibm.com

Volume Classification

When a data set is allocated, SMS follows a sequence of steps to decide on which volume or volume it should be placed. In order for a volume to be considered, it must be in one of the storage groups selected for the data set by the storage group ACS routine. SMS retrieves all the volumes in the selected storage groups and classifies them into four categories.

SMS attempts to spread initial allocations evenly across similar volumes and storage groups provided by the ACS storage group selection routine unless you specify Multi-Tiered SG (Y) in the storage class. In that case, SMS performs an allocation using storage groups in the order in which they are specified in the ACS storage group selection routine, and only those volumes that are in the first storage group listed are eligible for the primary list.

Volume selection for system-managed data sets first selects from the list of primary volumes. If there are no primary volumes, SMS selects from the secondary volumes. If there are no secondary volumes available, SMS selects from the tertiary candidates.

Primary volumes are those volumes which meet all the criteria specified in the storage class for the data set in addition to being online, below high threshold, and enabled at both volume and storage group level.

Secondary volumes are those which do not meet all of these criteria. For example, they may be quiesced or above their high thresholds.

Volumes are marked tertiary if the number of volumes in the storage group is less than the number of requested volumes (that is, the volume count specified by the user).
A common example of volumes being marked tertiary occurs when the tape mount management methodology is used. The original tape request might specify a volume count of 5. However, the allocation is redirected to a TMM buffer storage group containing only 2 volumes. In this case, the TMM buffer storage group would be marked tertiary, while a spill (quiesced) storage group containing 5 volumes is marked secondary. The result is that the quiesced volumes are preferred over the TMM volumes.

Volumes which do not meet the required specifications (for example, ACCESSIBILITY = CONTINUOUS or AVAILABILITY = STANDARD or CONTINUOUS) are marked rejected and will not be considered for selection. A detailed list of reasons why a volume may be rejected is included later in this presentation and can be found in information APAR II07464.

Data Set Separation

Data set separation allows you to designate groups of data sets in which all SMS-managed data sets within a group are kept separate, on the physical control unit (PCU) level, from all the other data sets in the same group.

To utilize data set separation, you must create a data set separation profile and specify the name of the profile to the base configuration. During allocation, SMS will attempt to separate the data sets listed in the profile.

A data set separation profile contains at least one data set separation group. Each data set separation group contains a fail level keyword value and a list of data set names to be separated from each other during allocation. The fail level keyword value specifies required or preferred data set separation.

Recommended Use of Separation

An existing data set might be allocated to one or more PCUs from which the current allocation is to be separated. In this case, SMS either rejects all volumes on the PCUs from the SMS volume preference list or places them at a less-preferred position. Volume rejection might drastically reduce the number of eligible volumes or result in allocation failures if SMS rejects all volumes.

Requesting data set separation can affect system performance and code path length. SMS must scan the profile to determine whether data set separation processing is needed for the current allocation. If the number of groups in the profile is large or if the number of names in the groups is large, this could take some processing time.

In addition, performance may be affected by the number of volumes that are defined to each physical control unit for which data set separation processing occurs.

Data set striping can use multiple controllers when multistriped. When using data set separation that references multistriped data sets, ensure that the storage group ACS routines select storage groups that contain a sufficient number of PCUs.

Generation data groups (GDGs) that are listed in the data set separation profile must specify the absolute generation and version number (for example, A.B.C.G0001V00). Data set separation does not support relative generation numbers and GDS base names. When specifying separation of GDGs, ensure that the number of generations that reside on pool volumes does not exceed the number of available PCUs.

IBM recommends that you migrate older GDG generations to allow a sufficient number of PCUs for the nonmigrated GDG generations.

Specifying Separation

A data set separation profile contains at least one data set separation group. Define the data set separation profile as a PDS member or sequential data set. Ensure that the profile is cataloged. Use a fixed record format or a fixed-block record format by specifying RECFM=FB. Then use the syntax shown to create a data set separation group.

You can specify either required or preferred data set separation by using the FAILLEVEL keyword. A FAILLEVEL of PCU indicates that separation on the PCU level is required. SMS fails the allocation if the requested data set cannot be separated on the PCU level from other data sets that are listed in the SEPARATIONGROUP. An allocation failure message is issued that indicates the number of volumes that were rejected by the data set separation profile.

A FAILLEVEL of NONE indicates that separation on the PCU level is preferred. SMS allows the allocation if the requested data set cannot be separated on the PCU level from other data sets that are listed in the SEPARATIONGROUP. SMS allocates the data sets on the same PCU and issues an allocation message that indicates that separation was not honored for a successful allocation. During volume selection, SMS treats all volumes on the same PCU level as less-preferred.

Data set names listed in a data set separation group must not contain quotes or wild cards, and must follow the naming convention described in the z/OS MVS JCL Reference. You can specify the same data set name in multiple data set separation groups.

The example of a data set separation profile shown contains two comments and two data set separation groups.

Multiple Separation Profiles

You can create multiple separation profiles in different data sets or PDS members, but only one separation profile can be active at any given time. This is because you can only specify one separation profile data set name in the SMS base configuration.

Alternatively, you can have a single separation profile which is used across multiple SMS configurations.

SMS reads the separation profile whenever it initializes and when a new configuration is activated.

When Separation Does Not Work

Under the environmental conditions listed, SMS might allocate on the same PCU two data sets that are specified for required or preferred separation.

Non-SMS-managed data sets cannot be used with separation because SMS does not control the placement of those data sets.

If the allocation is performed on a system in the SMS complex that runs a prior version of SMS that does not support data set separation, data set separation will not occur.

There are some services which use temporary data set names in allocating data sets and do not provide the real name of the data set to SMS.

If two data sets that are specified to be separated are allocated at the same time by two different tasks or two different systems, SMS may not be able to separate them. For performance reasons, SMS does not serialize data set allocations.

If the data set separation profile is modified after the configuration was activated, it can result in differences between the stored data set separation profile data set and the active copy of the data set separation profile.

Since rename processing does not involve physical movement of the data set, SMS does not perform separation during rename.

SMS does not perform separation processing during HSM migration to migration level one or two, but it does perform separation during HSM recall and recover.

SMS might not issue the corresponding warning message IGD17372I in those cases where allocation was successful and data set separation was not honored.

Conventional Volume Selection

Processing begins with a list of all the volumes in all of the candidate storage groups. Each volume is processed in turn and marked as primary, secondary, tertiary or rejected. This continues until all of the volumes in all of the candidate storage groups have been processed.

If a volume satisfies all of the primary conditions for allocation of the data set and the storage group it is in contains sufficient volumes to satisfy the allocation request, then the volume is marked primary. If the volume fails to meet all of the primary conditions, but the storage group it is in contains sufficient volumes to satisfy the allocation request, then it is may be marked secondary. If the storage group the volume is in does NOT contain sufficient volumes to satisfy the allocation request, then the volume may be marked tertiary.

Once all the volumes have been classified, SMS determines if there are any primary volumes. If there are, then it will always try to allocate the data set on those volumes first.

If there are no primary volumes or the allocation cannot be completed successfully on any of the primary volumes, then SMS determines if there are any secondary volumes. If there are, then SMS attempts the allocation on those volumes.

SMS only uses the tertiary list when there are no primary or secondary volumes or the allocation cannot be completed successfully from either of those lists.

Volume Selection Evaluation Process

The first and foremost consideration is data set separation. To be considered a primary candidate for this allocation, the volume must not reside behind a physical control unit that has allocated a data set from which the data set separation profile indicates this data set must be separated.

Next, SMS considers the volume count in the storage group. If there are enough volumes in the storage group to satisfy the requested volume count, then SMS considers the storage group. If not, then all the volumes in the storage group are placed on the tertiary list.

The next weighting factor is the high threshold. To be considered a primary candidate, a volume must have sufficient space for the allocation amount without exceeding its storage group's high threshold value.

After that, the SMS status of the volume and storage group is considered. To be primary, both the storage group and the volume must be enabled. If either or both is quiesced, then the volume becomes a secondary candidate. If either is disabled, then the volume is rejected.

For end of volume extends, only the primary storage group (that is, the one where it was initially allocated) and extend storage group are eligible for selection. SMS prefers volumes that are in the primary storage group for the data set over those volumes in the extend storage group specified for the primary storage group.

For initial allocation, a non-overflow storage group is always preferred over an overflow storage group.

Next, SMS considers the storage class attributes. First is the initial access response time. A non-zero IART value indicates to SMS that mountable volumes are to be used, if there are any. SMS will only allocate a data set with a non-zero IART value on DASD if there are no mountable volumes which can be used.

VTS (Virtual Tape Server) Cache Management: An initial access response time (IART) of 100 or greater means the volume has least preference in the cache. 0-99 means the volume has most preference.

Object Access: You can use the Initial Access Response Seconds field to specify the desired response time (in seconds) required for locating, mounting, and preparing a piece of media for data transfer. OAM uses this value to interpret the storage level, that is, to place an object at an

appropriate level in the object storage hierarchy. For objects, both the IART and the sustained data rate (SDR) are applicable.

OAM uses IART as follows:

If the IART value is 0, OAM writes to DASD.

If the IART value is 1-9999, OAM selects removable media, either optical or tape. If the SDR is greater than or equal to 3, the primary copy of the object is stored on a tape volume. If the SDR for the object is less than 3, the primary copy of the object is stored on an optical disk volume.

The IART for an optical volume depends on its location. The time required for an operator to mount a shelf-resident optical volume is significantly longer than that for the automatic mounting of a library-resident optical volume. You can use up to four characters to specify a valid value of 0-9999, or leave the field blank. The default is blank and selects fixed DASD. However, an OAM request fails if it tries to use a storage class with a blank IART value.

Data Set Access: SMS allows the system resources manager (SRM) to select a DASD volume from the primary volume list if the IART value is 0 or unspecified. SRM volume selection is ideally suited for batch jobs.

You can allow SMS, instead of SRM, to select a DASD volume by assigning a storage class with a non-zero IART value. In this case, none of the eligible volumes is placed on the primary list; all of the eligible volumes are placed only on the secondary or tertiary list. SMS uses the randomizing technique, for example, so as not to always select the volume that has the most free space.

After that, SMS looks at the ACCESSIBILITY specification. The storage class accessibility attribute defines the function of the hardware supporting the point-in-time copy, using either concurrent copy, virtual concurrent copy, or FlashCopy. The point-in-time copy allows database management systems (DBMSs) to take what appears to be an instantaneous copy of data or a "fast" point-in-time copy. The copy can be for backup purposes (generally to tape) or for copying a database from one set of DASD volumes to another. The accessibility attribute allows you to direct allocation of new data sets to DASDs connected to an IBM 3990 Storage Control unit with cache that supports the concurrent copy function, IBM RAMAC Virtual Array (RVA) devices with the SnapShot copy support, DFSMSdss with the virtual concurrent copy, and the IBM ESS with the FlashCopy service.

The IBM Enterprise Storage Server (ESS) FlashCopy service provides the appearance of instantaneous replication of a range of track images. You can invoke this service with DFSMSdss and use it to create copies for disaster recoveries, business intelligence applications, data in a test environment, and instantaneous checkpoints.

More information on the ESS FlashCopy service can be found in "DFSMS Advanced Copy Services" and "DFSMS/MVS Support for the IBM Enterprise Storage Server".

When you specify accessibility attributes, you identify whether you want SMS to use versioning or backup devices for either concurrent copy or virtual concurrent copy.

Versioning: Device can create a "fast" point-in-time version of a data set, which is then available for application testing, reporting, or backup operations. While the version is being made, the data set is unavailable for normal application processing for a minimal period of time. Versioning is done using the SnapShot feature of the RAMAC Virtual Array or the ESS FlashCopy service.

Backup: Device can create a "fast" point-in-time backup copy of a data set. While the backup copy is being made, the data set is unavailable for normal application processing for a minimal period of time. Two methods are supported:

Method 1: Establish a concurrent copy session with the 3990 DASD controller and make the backup copy.

Method 2: Use virtual concurrent copy. DFSMSdss uses the SnapShot feature of the RAMAC Virtual Array or the ESS FlashCopy service to create a point-in-time copy and then does I/O to create the backup to whatever target device you specified.

You can use a combination of values for the accessibility attribute and its subparameters, Versioning and Backup, to request the point-in-time copy. In the Accessibility field, you specify whether allocation to a point-in-time copy-capable volume is required (CONTINUOUS), preferred (CONTINUOUS PREFERRED), or discouraged (STANDARD). You then specify values for the Versioning and Backup subparameters to select which devices you want used for the copy.

The following defines your allocation request to a point-in-time copy-capable volume:

CONTINUOUS (C): Only point-in-time copy volumes are selected.

CONTINUOUS PREFERRED (P): Point-in-time copy volumes are preferred over non-point-in-time copy volumes.

STANDARD (S): Non-point-in-time copy volumes are preferred over point-in-time copy volumes.

NOPREF (N): Point-in-time copy capability is ignored during volume selection. This is the default. The first and foremost consideration is data set separation. To be considered a primary candidate for this allocation, the volume must not reside behind a physical control unit that has allocated a data set from which the data set separation profile indicates this data set must be separated.

Having done that, SMS checks the storage class availability attribute. The Availability field is used to specify whether data set access should continue in the event of a single device failure. Storage classes with a blank Availability field default to NOPREF.

CONTINUOUS: Specify an availability of CONTINUOUS if you do not want a device failure to affect processing. Only duplexed and RAID volumes are eligible for this setting.

If CONTINUOUS availability is specified, data is placed on a device that can guarantee that it can still access the data in the event of a single device failure. This option can be met by

A dual copy volume
An array DASD

PREFERRED: Array DASD volumes are preferred over non-duplexed volumes. Dual Copy volumes are not candidates for selection.

STANDARD: If data sets do not require such a high level of availability, specify STANDARD availability, which represents normal storage needs.

Specify an availability of STANDARD to cause processing of a data set to stop after a device failure. Simplex volumes are preferred over array DASD. SMS selects only volumes that are not dual copy. This attribute does not apply to objects. Array DASD are acceptable candidates for both STANDARD and CONTINUOUS availability requests.

Prior to DFSMS/MVS V1R2, the default of STANDARD acts the same as the default of NOPREF. Change STANDARD to
NOPREF for compatibility with releases prior to DFSMS/MVS V1R2.

NOPREF: Simplex and array DASD are equally considered for volume selection. NOPREF is the default. Dual copy volumes are not candidates for selection.

Note: On DFSMS/MVS V1R2, SMS interprets a blank Availability field as STANDARD instead of NOPREF. To ensure that storage classes are set correctly on a DFSMS/MVS V1R2 system, issue ISMF ALTER and SAVE commands against those storage classes which have a blank Availability field or which specify NOPREF. You do not need to modify the storage classes.

Next, SMS looks at the data class. If the data class indicates that the data set must be allocated in extended format, then a volume which supports this must be available. If it is not, then the allocation fails. If the data class indicates that extended format is preferred, but not required, then SMS attempts to allocate the data set on a volume which supports extended format, but it will allocate the data set on a non-extended format device if no such device is available.

The MSR serves two purposes in SMS. First, it is used as the performance objective for selecting candidate volumes for new data set placement. During a new data set allocation, SMS searches for a volume that meets or closely matches this objective. If no volume satisfies the objective, then SMS attempts to find a volume that comes closest to matching it. If more than one MSR is

explicitly or implicitly specified, the storage class and associated device MSRs are averaged and compared.

Second, if the data is placed on a volume attached through an IBM 3990 Storage Control with cache, and cache is enabled for that volume, the MSR is used to determine if caching is mandatory, optional, or should be inhibited for the data set. This attribute does not apply to objects.

You can request SMS to ignore various device performances during volume selection by leaving all MSR and BIAS fields blank. This lets you spread data evenly across non-cached and cache active devices.

Millisecond Response Time and Data Set Allocation: DASD can have different performance capabilities for direct access (random access, for example) and for sequential access applications. Its performance capabilities depend on whether you are reading data or writing data.

Each device type and model has a predetermined MSR capability for each condition. Additionally, if the device is attached to a cache capable control unit, the response capabilities are improved when caching is active. Therefore, each device is represented by eight MSR values:

Uncached Performance
- Direct Read MSR
- Direct Write MSR
- Sequential Read MSR
- Sequential Write MSR

Cache performance, DASD Fast Write performance, or both (if active)
- Direct Read MSR
- Direct Write MSR
- Sequential Read MSR
- Sequential Write MSR

If a device is cache capable, it must also have caching active at the time of allocation in order to be represented by the caching MSR values.

More information on this can be found in "OS/390 V2R10.0 DFSMSdfp Storage Administration Reference."

After the system selects the primary space allocation volume, that volume's associated storage group is used to select any remaining volumes requested for the data set.

Any criterion that is not requested or not met gets a preference value of 0.

Example

Since volume B has a higher total preference value, SMS selects volume B. If there had been a volume C which had the same total preference value as B, then SMS would randomly select either B or C.

The Primary List

All volumes in all the specified storage groups are candidates for the first, or primary, list. The primary list consists of those devices that fit all the storage class and data class requirements, are online, enabled and below threshold. All devices on this list are considered equally qualified to satisfy the request.

Whenever possible, a choice is made from among the candidates on this list, essentially choosing the device that is the least busy (that is, by preferring volumes that are not already allocated to a job or subsystem and that have the least I/O delay as determine by SRM). It is possible that there are no volumes on the primary list if no devices in the storage group(s) meet all of the primary list criteria. It is also possible that even though there are devices on the primary list, the data set cannot be allocated successfully to any of these devices (i.e. there is not enough space available when the actual allocation request is made.)

For example, if you specify an MSR of 25 is your storage class, volumes that are close to this MSR would be POOL1 volumes. POOL1 volumes are those volumes that meet your MSR requirement by a certain percentage. Each MSR percentage represents a range of MSRs. All volumes that fall within that MSR percentage are considered equal in performance. Volumes that provide a faster MSR by a larger percentage than POOL1 volumes are considered POOL2 volumes. POOL2 volumes include POOL1 volumes plus the volumes in the next higher MSR percentage. As shown in the earlier diagram, POOL1 volumes are preferred over POOL2 volumes.

Once a list of primary volumes is selected, it is passed to the system resource manager (SRM) which selects a volume based on channel resources.

Volume selection from the primary list can result in skewing in extreme circumstances. These circumstances include the times when a VTOC index is disabled or when a new volume is added to SMS and selected most of the time. Starting with DFSMS/MVS V1R3, SMS uses the randomizing technique not to favor one volume over others, such as an empty volume, a volume with a disabled index, or a volume under the allocation threshold. You can force SMS to always use the randomizing technique by specifying a non-zero IART value.

The Secondary List

The secondary list contains all the volumes that can be used for the data set allocation, even if they don't meet all the storage class criteria. It also includes all primary volumes.

If allocation is not successful from the primary list, then it is attempted from this list, starting at the top and working down to the bottom until successful allocation is achieved. The order of the secondary list is established to select the preference order listed on the slide.

In order to prevent X37 ABENDs, the first volumes on the list are those that can allocate the data set while remaining below their high thresholds. Second, enabled volumes are preferred over volumes which are quiesced. Third, SMS attempts to meet the volume characteristics requested by the data class and storage class as closely as possible. And lastly, SMS considers the MSR requested to match it to the performance characteristics of the volume.

SMS does perform some randomization of the list in order to prevent it from selecting the same volume over and over.

The Tertiary List

When a storage group does not contain enough volumes to satisfy the volume count, all volumes in the storage group are flagged as tertiary. Tertiary volumes are only selected when there are no primary or secondary volumes and the request is for a non-VSAM Non-guaranteed SPACE request. The concept of tertiary volumes does not apply to VSAM data sets. In other words, for all VSAM Non-guaranteed SPACE requests, the volume count does not play a role in determining which storage group is selected.

Why Isn't My Volume Primary?

1. A detailed list of reasons for volume rejection is provided on a subsequent foil.
2. The amount of space requested would have pushed the volume over its high threshold.
3. Quiesced volumes and storage groups are never considered primary.
4. A volume which is not close to the requested MSR is not considered primary. Volumes which can meet the MSR are preferred. APAR OW08630 allows provides the options of specifying that there is no performance preference for SMS volume selection.
5. If the VTOC index is broken or disabled, SMS is not updated with space statistics by CVAF. This may also happen if OEM products bypass CVAF, the component that notifies SMS of space changes. SMS not being updated with space statistics can also result in a volume being overused.
6. A non-zero initial access response time value in the storage class can result in all DASD devices being removed from the primary list.

VTS (Virtual Tape Server) Cache Management: An initial access response time (IART) of 100 or greater means the volume has least preference in the cache. 0-99 means the volume has most preference.

Object Access: You can use the Initial Access Response Seconds field to pecify the desired response time (in seconds) required for locating, mounting, and preparing a piece of media for data transfer. OAM uses this value to interpret the storage level, that is, to place an object at an appropriate level in the object storage hierarchy. For objects, both the IART and the sustained data rate (SDR) are applicable.

OAM uses IART as follows:

If the IART value is 0, OAM writes to DASD.
If the IART value is 1-9999, OAM selects removable media, either optical or tape. If the SDR is greater than or equal to 3, the primary copy of the object is stored on a tape volume. If the SDR for the object is less than 3, the primary copy of the object is stored on an optical disk volume.

The IART for an optical volume depends on its location. The time required for an operator to mount a shelf-resident optical volume is significantly longer than that for the automatic mounting of a library-resident optical volume. You can use up to four characters to specify a valid value of 0-9999, or leave the field blank. The default is blank and selects fixed DASD. However, an OAM request fails if it tries to use a storage class with a blank IART value.
Data Set Access: SMS allows the system resources manager (SRM) to select a DASD volume from the primary volume list if the IART value is 0 or unspecified. SRM volume selection is ideally suited for batch jobs.

You can allow SMS, instead of SRM, to select a DASD volume by assigning a storage class with a nonzero IART value. In this case, none of the eligible volumes is placed on the primary list; all of the eligible volumes are placed only on the secondary or tertiary list. SMS uses the randomizing technique, for example, so as not to always select the volume that has the most free space.

7. If the controller is IMLed while online to MVS, the MVS device control blocks may not reflect the current state of the volume. The device(s) should be varied off-line and then back online to update the MVS control block status.
8. Accessibility=Preferred in the storage class causes all volumes without 3390 extended platform support to be removed from the primary list. Accessibility=Standard in the storage class causes all volumes with 3990 extended platform support to be removed from the primary list. Likewise, accessibility=preferred causes only volumes with 3990 extended platform support to be placed in the primary list.
9. Availability=Standard causes all devices with 3990 extended platform support to be excluded from the primary list. If allocations should be evenly spread between RAID and simplex devices, then Availability=No Preference should be used. If Availability=Preferred is specified, simplex devices are removed from the primary list.
10. When extended format is specified as preferred in the data class, only volumes which are behind a controller which supports extended format are considered primary.
11. If the number of volumes in the storage group is less than the requested volume count, then all volume in the storage group are placed on the tertiary list, even if the number of volumes

defined to the storage group is sufficient, but not enough are currently available.

For example, this happens when a tape request of five volumes causes a tape mount management buffer storage group containing only two volumes to be marked tertiary, while an overflow storage group or a quiesced storage group containing the requested five volumes is marked secondary. The overflow volumes and quiesced volumes on the secondary list are preferred over the tape mount management volumes on the tertiary list. SMS will attempt to allocate space using the overflow volumes first, then the quiesced volumes, and lastly the tape mount management buffer volumes.

Why Was My Volume Rejected?

1. In order for the volume to be used, it must be in one of the storage groups selected for the data set by the storage group ACS routine.
2. Volumes which are not online cannot be used.
3. Volumes must be enabled or quiesced (QUIALL, QUINEW); they cannot be disabled (DISNEW, DISALL) or not connected (NOTCON).
4. Volumes must have enough space to satisfy the primary space quantity requested.
5. The VTOC must have sufficient space to add the new information about the data set.
6. The availability requirement is determined by the storage class specification. Continuous availability requires dual copy or array DASD. Standard availability requires non-dual copy or array DASD.
7. If Accessibility=Continuous, then the data set must be allocated on a volume which supports concurrent copy or on Array DASD. If one cannot be found, the allocation fails.
8. It is not sufficient simply to add a volume to the SMS configuration. It must also be initialized with the attributes that mark it as being SMS-managed.
9. If DSNTYPE=Extended,Required, the data set must be allocated on a volume which supports extended format. A volume supports extended format only when it is behind a controller which supports extended format.
10. Initial access response time is specified in the storage class. When it is zero or blank on DFSMS/MVS 1.2.0 and above, a fixed DASD volume must be selected (as opposed to an optical device; D/T3995-151 or D/T3995-153).
11. A volume is placed in the exclude list if (a) The data set is being extended to a new volume, and the volume is one of the volumes on which the data set already resides. (b) Allocation was attempted on the volume and failed. (c) The volume does not meet one of the requirements for the data set such as availability or accessibility. (d) The MVS unit eligibility service rejected the volume.
12. A volume is placed in the include list if: (a) This is a guaranteed space request and the volume was explicitly specified. (b) This is a recall or recover for a guaranteed space data set, and DFSMShsm or DFSMSdss passed the volume(s) on which the data set originally resided as input. (c) The data set is being extended to a new volume, does not already reside on the volume, and the volume is in the storage group in which the data set resides. (d) The volume is an eligible candidate for initial allocation using the conditions previously discussed.
13. The UCB type is considered unusable if the UCB pointer for the volume is zero, the pointer to the device class extension (DCE) from the UCB is zero, the pointer to the device performance characteristics table (DPCT) from the DCE is zero, or the DCE valid features bit (DCEVFEAT) is off.

14. If guaranteed space is used, the user may either explicitly specify the volumes to be used or allow the system to choose the volumes. If the user explicitly specifies volume serials and the number of volume serials matches the volume count for the data set, then only those volumes may be used.
15. If the storage class for a data set specifies guaranteed space, then any volume which is in a storage group with insufficient volumes to satisfy the volume count for the request is rejected. This includes the case where the number of volumes defined to the storage group is sufficient, but not enough of the volumes are currently available.

Note: With the IBM Enterprise Storage Server (ESS), the guaranteed space attribute of a storage class with specific volume serials is no longer required for data sets other than those that need to be separated (for example, DB2 online logs and BSDS) or that must reside on specific volumes because of their naming conventions (for example, VSAM RLS control data sets). The ESS storage controllers use the RAID architecture that enables multiple logical volumes to be mapped on a single physical RAID group. If required, you can still separate data sets on a physical controller boundary for the purpose of availability.

The ESS capabilities of multiple allegiance and parallel access volumes (PAV), along with its bandwidth and caching algorithms, make it unnecessary to separate data sets for the purpose of performance. Traditionally, IBM storage subsystems allow only one channel program to be active on a disk volume at a time. This means that once the subsystem accepts an I/O request for a particular unit address, this unit address appears "busy" to subsequent I/O requests. This ensures that additional requesting channel programs cannot alter data that is already being accessed. By contrast, the ESS is capable of multiple allegiance, or concurrent execution of multiple requests from multiple hosts. That is, the ESS can queue and concurrently execute multiple requests for the same unit address from multiple hosts, provided that no extent conflict occurs.

The ESS can also execute multiple concurrent accesses to a single volume from a single host or PAV. To access a volume concurrently, you must associate multiple device numbers with a single volume. The ESS provides this capability by allowing you to define a PAV-base address and one or more PAV-alias addresses. It allows up to 255 aliases per logical volume. Therefore, you no longer have to separate data sets from each other for performance reasons.
16. SMS called DADSM to allocate space on the volume and DADSM returned an unsuccessful return and reason code.
17. If a volume is sufficiently fragmented, it may not be possible to find enough contiguous space to allocate the primary space extent.
18. During extend (EOV) processing, volume selection rejects volumes whose device types do not match the device type of the volume that the data set currently resides on.

19. For multistripe data sets, all volumes in a storage group which contains mixed sustained data rates (for example, a mixture of 3380s and 3390s) are rejected.

20. IMLing a controller while online to MVS may result in the MVS device control blocks not reflecting the current state of the volume, leading SMS to reject the volume. For performance reasons and to avoid excessive I/O, SMS uses the status in the MVS device control blocks. ISMF displays correct device statuses since it queries the device/controller directly. The device(s) should be varied off-line and back online to update the MVS control block status.

Wrong Volume Selected?

It may appear that the wrong volume is being selected if your ACS routines are not doing what you expected. Small logic errors can result in unpleasant surprises. The simplest form of this is when there is an error in the storage group ACS routine which causes a storage group other than the one(s) expected to be assigned. However, if your storage group ACS routine keys off the storage class, management class, or data class assigned, then a mistake in one of the other routines could have this effect as well.

Another factor which affects volume selection is channel path
utilization. SMS passes all the volumes placed on the primary list are to SRM. SRM then selects volumes based on channel path utilization.

It could be the case that the volume selected was on the primary list and the volume you expected SMS to use was placed on the secondary list. The list of reasons why this can occur is detailed on the next foil.

If there are no primary volumes, or attempts to allocate the data set on all of the primary volumes fail, then secondary or tertiary candidates could be used.

Effect of MSR and Bias

The MSR used from the device performance characteristics table (DPCT) for a device is the current capability of the device. If the device currently has cache active, then the device's MSR is for the device with cache active. If the device currently has cache turned off, then the device's MSR will reflect the native capability of the device. The MSR is the only consideration used when determining whether or not a device is a primary candidate. If it provides an MSR close to the requested MSR, then it is place on the primary list.

Be careful when tuning your storage class ACS routines to match the mixture of volume capabilities assigned by your storage group ACS routines. For example, if you have both cached and noncached 3390-3 in the assigned storage groups and your ACS routine assigned the same storage class (MSR=25) to all data sets, the noncached volumes are preferred over the cached volumes. This results in the noncached volumes being overused and the cache-active volumes are not selected until all the noncached volumes are above high threshold. If an MSR=2 is specified, then the cached volumes are overused.

This also holds true for other attributes. If you always assign ACCESSIBILITY=PREFERRED, then concurrent copy volumes are overused.

The MSR and bias values you specify in the storage class can be used to determine how buffers are to be allocated when system-managed buffering is used for VSAM applications. To do this, specify a value of System for the Record Access Bias attribute in the data class. This capability applies only to system-managed VSAM data sets allocated in extended format and accessed by batch applications. MSR values are also used to cache PDSE members and HFS files.

Effect of MSR and SDR

Coding an MSR of 999 in the storage class causes devices with a native value (device performance without the use of cache or DFW) of 25 to be placed on the primary list.

Sustained Data Rate applies to physical sequential data sets, VSAM data sets, and objects. For Physical Sequential and VSAM data sets the Sustained Data Rate has an effect only for extended format data sets. Striping allows you to spread data across different DASD volumes and controllers. The number of stripes is the number of volumes on which the data set is initially allocated. Striped data sets must be system-managed and must be in an extended format. When no volumes that use striping are available, the data set is allocated as nonstriped with extended preferred specified in the data class; the allocation fails if extended required is specified in the data class.

Physical sequential (PS) data sets cannot be extended if any of the stripes cannot be extended. For VSAM data sets, each stripe can be extended to an available candidate volume if extensions fail on the current volume.

If you do not specify Guaranteed Space=Y in the storage class, the following is true:

- The system uses the SDR field in conjunction with the assumed device transfer rate to derive the number of stripes that a striped data set can have. The SDR specifies the target throughput rate. Valid values are blank or from 0 to 999. The maximum number of stripes for PS data sets is 59. For VSAM data sets, the maximum number of stripes is 16.
- For nonguaranteed space requests, if the SDR field is blank or 0, the target number of stripes is 1. If the value in the SDR field is greater than 1, it is divided by 4 for 3390s, or by 3 for 3380s to determine the stripe count. For example, for an SDR value of 24, a storage group of 3380s would have a stripe count of 8, and a storage group of 3390s would have a stripe count of 6. The volume must be able to satisfy the requested primary space.

- Depending on the number of eligible volumes in the candidate storage groups, the actual number of stripes allocated to a data set can be less than the derived one. As the stripe count is decreased, the primary space for each stripe increases, except for the guaranteed space requests.

If a sustained data rate (SDR) greater than zero is specified, conventional volume selection is not used, and MSR, Availability, Accessibility, and free space criteria may be ignored when selecting a volume. An SDR greater than zero causes striping volume selection to be used; this method of volume selection is more concerned with controllers than volume attributes.

Object Access: OAM uses the SDR field to determine whether the primary copy of the object should be placed on optical or tape media. An SDR value of 3 or greater places the data on tape media, and a value of 2 or less places the data on optical media.

Effect of Availability

If continuous is specified, the allocation must occur on Dual Copy or Array DASD. Note that dual copy volumes are only eligible if availability is continuous. If no dual copy or array DASD volumes are available, the allocation fails.

If preferred is specified, Array DASD volumes are preferred over simplex volumes, but if none are available, the allocation will be attempted on simplex volumes.

If standard is specified, simplex volumes are preferred over Array DASD, but array DASD will be used if no simplex volume is available.

If no preference is specified, simplex and array DASD are considered equal for volume selection. This is the default.

When preferred is specified, the primary list consists of array DASD. The secondary and tertiary lists will prefer array DASD over simplex volumes.

When standard is specified, the primary list consists of simplex volumes. In the secondary and tertiary lists, array DASD and simplex volumes are treated equally. Dual copy volumes are not eligible.

Effect of Accessibility

The storage class accessibility attribute defines the function of the hardware supporting the point-in-time copy, using either concurrent copy, virtual concurrent copy, or FlashCopy. The point-in-time copy allows database management systems (DBMSs) to take what appears to be an instantaneous copy of data or a "fast" point-in-time copy. The copy can be for backup purposes (generally to tape) or for copying a database from one set of DASD volumes to another. The accessibility attribute allows you to direct allocation of new data sets to DASDs connected to an

IBM 3990 Storage Control unit with cache that supports the concurrent copy

function, IBM RAMAC Virtual Array (RVA) devices with the SnapShot copy support, DFSMSdss with the virtual concurrent copy, and the IBM ESS with the FlashCopy service.

The IBM Enterprise Storage Server (ESS) FlashCopy service provides the appearance of instantaneous replication of a range of track images. You can invoke this service with DFSMSdss and use it to create copies for disaster recoveries, business intelligence applications, data in a test environment, and instantaneous checkpoints.

See z/OS DFSMS Advanced Copy Services and z/OS DFSMS Software Support for the IBM Enterprise Storage Server for further information on the ESS FlashCopy service.

When you specify accessibility attributes, you identify whether you want SMS to use versioning or backup devices for either concurrent copy or virtual concurrent copy.

Versioning Device can create a "fast" point-in-time version of a data set, which is then available for application testing, reporting, or backup operations. While the version is being made, the data set is unavailable for normal application processing for a minimal period of time. Versioning is done using the SnapShot feature of the RAMAC Virtual Array or the ESS FlashCopy service.

Backup Device can create a "fast" point-in-time backup copy of a data set. While the backup copy is being made, the data set is unavailable for normal application processing for a minimal period of time. Two methods are supported:

Method 1 Establish a concurrent copy session with the 3990 DASD controller and make the backup copy.

Method 2 Use virtual concurrent copy. DFSMSdss uses the SnapShot feature of the RAMAC Virtual Array or the ESS FlashCopy service to create a point-in-time copy and then does I/O to create the backup to whatever target device you specified.

You can use a combination of values for the accessibility attribute and its subparameters, Versioning and Backup, to request the point-in-time copy. In the Accessibility field, you specify whether allocation to a point-in-time copy-capable volume is required (CONTINUOUS), preferred (CONTINUOUS PREFERRED), or discouraged (STANDARD). You then specify values for the Versioning and Backup subparameters to select which devices you want used for the copy.

The following defines your allocation request to a point-in-time copy-capable volume:

CONTINUOUS (C) Only point-in-time copy volumes are selected.

CONTINUOUS PREFERRED (P) Point-in-time copy volumes are preferred over non-point-in-time copy volumes.

STANDARD (S) Non-point-in-time copy volumes are preferred over point-in-time copy volumes.

NOPREF (N) Point-in-time copy capability is ignored during volume selection. This is the default.

This chart and the one that follows it show how devices are selected based on the accessibility option and its subparameters. By careful specification of these value, you can specify that only versioning devices are to be used, only backup devices are to be used, or that any device may be used.

Multi-Tiered Storage Groups

You can specify Multi-Tiered SG (Y) in the storage class to enable the multi-tiered storage group function. When you specify this function, SMS allocates to enabled pool storage groups in the order they are specified in the ACS storage group selection routines.

For example, if Multi-Tiered SG (Y) is specified in the storage class and an ACS storage group routine assigns the following three similar storage groups:

```
SET &STORGRP = 'SG1', 'SG2', 'SG3'
```

The result is that SMS selects the enabled volumes that are below the high threshold in SG1 before selecting volumes in the SG2 storage groups that are listed second, third, and so forth, can also be used depending on the status of the underlying volumes when:

- All SG1 enabled volumes exceed high threshold, then SMS selects the enabled volumes in SG2.
- All SG1 and SG2 enabled volumes exceed the high threshold, SMS selects the enabled volumes in SG3.
- All enabled volumes exceed the high threshold, SMS selects the quiesced volumes in the same storage group order.

Only volumes that reside in the first storage group assigned by the ACS routines (SG1 in the example) are eligible for the primary list.

Restrictions: The multi-tiered storage group attribute is only used with conventional volume selection. This attribute is ignored if Striping Volume Selection is used.

Parallel Access Volumes

If the Enterprise Storage Server (ESS) parallel access volume (PAV) option is enabled, you can use the DFSMS storage class parallel access volume (PAV) option to ensure taht data sets that require high performance are only allocated to volumes on which the ESS PAV option is enabled.

The DFSMS PAV capability option includes the following settings with these results:
Required Only volumes with the PAV feature enabled are selected.

Preferred Volumes with the PAV feature enabled are eligible to be primary volumes. Volumes without the PAV feature enabled are only eligible to be secondary volumes.

Standard Volumes without the PAV feature enabled are preferred over volumes with the PAV feature enabled and are eligible to be primary volumes. Volumes with the PAV feature enabled are only eligible to be secondary volumes.

Nopreference Whether the PAV feature is enabled or not for a volume is ignored and has no effect on the volume selection process. This is the default value for this option.

Secondary/Tertiary List Preferencing

1. If data set separation is requested, then the first criteria used to select a volume for the data set is that it be placed on a different PCU from the data sets from which it is being separated.

2. The next consideration is to try to meet the requested volume count. If the storage group cannot provide enough volumes to meet the count, then it will be less preferred than one which can.

3. A volume below high threshold is preferred over a volume which is above high threshold. A volume is considered below high threshold when the primary space allocation would not cause the high threshold limit to be exceeded.

4. Within groups of volumes which are below and above threshold, volumes are divided based on status. Enabled volumes are preferred over quiesced volumes. A volume is considered quiesced if either it or its storage group is quiesced. The order of preference is:
a. Below threshold and enabled
b. Below threshold and quiesced
c. Above threshold and enabled
d. Above threshold and quiesced

5. An initial access response time which is greater than zero causes mountable volumes (optical) to be preferred over fixed DASD. The order of preference is:
a. Below threshold, enabled, mountable
b. Below threshold, enabled, fixed DASD
c. Below threshold, quiesced, mountable

d. Below threshold, quiesced, fixed DASD

e. Above threshold, enabled, mountable
f. Above threshold, enabled, fixed DASD
g. Above threshold, quiesced, mountable
h. Above threshold, quiesced, fixed DASD

6. If Accessibility=Preferred volumes which support concurrent copy, SnapShot or FlashCopy are preferred over volumes which do not. (All volumes behind a 3990-6 or a 3990-3 controller with extended platform microcode are considered to support concurrent copy.) The NOPREF value can be entered for accessibility and allows extended platform and non-extended platform volumes to be treated equally (included in the primary list). For a multivolume data set, if any extent resides on a non-extended platform device, the entire data set is not accessible for concurrent copy.

7. Extended Format=Preferred specified in the data class causes volumes behind a controller that supports extended format to be preferred over volumes which do not.

8.Availability=Preferred causes array DASD volumes to be placed ahead of non-duplexed volumes.

When RAMAC devices are placed in dual copy mode, SMS volume selection sees the devices as dual copy, not RAMAC. RAMAC devices are eligible for all availability requests unless they are placed in dual copy mode. Dual copy devices can only be used for Availability=Continuous. For any other availability specification, only non-dual copy devices can be used.

9. Care should be taken to tune your storage classes and storage class ACS routine to match the device capabilities of your storage groups. If you have both cached and non-cached devices in the assigned storage groups and your ACS routine assigns the same storage class (MSR=25) to all data sets, the non-cached devices are preferred over the cached devices. The non-cached devices are over utilized, and the cached devices are not selected until the non-cached devices are above high threshold. If MSR=2 is specified, then the reverse occurs. Recommended MSR settings can be found in DFSMSdfp Storage Administration Reference (SC26-4920) under "Understanding Millisecond Response Time" and in documentation APAR OW08472.

The same caution applies to other attributes as well. If an accessibility of preferred is always used, then concurrent copy (T0) capable volumes will be over utilized. In general, if a storage class is assigned to data sets which occupy 30% of the total space of all selected storage groups, then the selected storage groups (when combined) should have 30% of their space backed by volumes which match all the storage class parameters.

Take care when specifying multiple preferences (AVAILABILITY=P, ACCESSIBILITY=P, MSR/BIAS value specified, If EXT=P, IART>0). Each additional preference specified adds additional work to SMS volume selection routines and might affect allocation performance.

10. MSR=999 indicates that the data set should never be cached. As a result, all non-cache active volumes are placed ahead of cache active volumes.

Spreading Allocations Across Volumes

You might want to separate allocations and spread them across multiple volumes to reduce I/O contention and improve performance. After analyzing your data, you might find that certain data may not need some of the advanced storage features in your environment.

If you want to spread conventional allocations across volumes of different features, use one or all of the attribute settings suggested here.

Overflow Storage Groups

Overflow storage groups are designated to handle periods of high demand for initial primary space allocations. They are utilized when the enabled (non-overflow) storage groups are above high threshold. During initial allocation, volumes in an overflow storage group are less preferred than volumes in an enabled storage group but more preferred than volumes in a quiesced storage group. Therefore, overflow volumes are not placed on the primary volume list but can be placed on the secondary volume list. When an overflow storage group contains more volumes than a buffer storage group, specified volume counts might result in overflow volumes being preferred over volumes in the buffer storage group.

Volumes residing in overflow storage groups are preferred over quiesced volumes and storage groups. If you quiesce an overflow storage group or volume then the quiesced volumes are preferred over quiesced overflow volumes.

Because overflow storage group status is not recognized on systems prior to z/OS R1V3, IBM recommends setting the overflow storage group status to "quiesced" on lower-level systems. This preserves overflow storage groups for peak usage.

Extend Storage Groups

An extend storage group is a pool storage group to which data sets from the primary storage group can be extended when there is an insufficient amount of storage on the primary storage group. A primary storage group is the storage group in which the initial allocation resides.

For each primary storage groups, you can define only one extend storage group. However, you can define the same extend storage group to more than one primary storage group. Also, you can define two storage groups as extend storage groups of each other.

Extend storage groups will not be used for initial allocation unless they are specified in the ACS routines. All storage groups that are listed in the ACS routines are candidates for initial allocation. Extend storage group attributes are not referenced during initial allocation.

Example: You can define storage group (SG) 3 as an extend of SG 2, and define SG 2 as an extend of SG 1. If SMS selects SG 1 as the primary storage group, the data set can extend only to SG 2. If SMS selects SG 2 as the primary storage group, the data set can extend only to SG 3.

When you specify an extend storage group, you must ensure that connectivity across the SMSplex is the same for both the primary storage group and the extend storage group.

A storage group may be both an overflow storage group and an extend storage group.

Striping Volume Selection

Volumes are classified as primary and secondary, and primary volumes are preferred over secondary volumes. A single primary volume is randomly selected for each unique controller. All other eligible volumes behind the same controller are secondary. Secondary volumes are randomly selected if initial allocation on the primary volume is unsuccessful. There is no preference in different controller models as long as the controller supports striping.

SMS randomly selects a storage group from a list of storage groups that contain as many unique controllers as the requested stripe count. If there are no storage groups with enough unique controllers to meet the requested stripe count, the storage group with the largest number of unique controllers is selected (randomly, if there are more than one).

Volumes that meet the requested MSR are preferred over volumes that do not. A volume is considered to meet the requested MSR if its performance falls within a predetermined range of the requested MSR.

The throughput of striped data sets is gated by the slowest device if the striped set includes devices of varying data delivery capabilities.

Storage groups containing mixed device types (and thus mixed sustained data rates) are not considered.

Eligible volumes must remain below the high threshold defined in the storage group after the stripe is allocated. They must be able to satisfy the primary space requested by the number of stripes. This is different from non-striped allocations where no volume is rejected because of space, although they can be given a lower priority.

Quiesced and enabled volumes are treated identically for striped allocations, unlike non-striped allocations where

enabled volumes are preferred over quiesced volumes.

The number of target volumes is computed by dividing the SDR specified in the storage class by a value of 3 for 3380 devices, and by a value of 4 for 3390 devices, and rounding up the result if required. For example, an SDR of 18 results in a target stripe count of 6 on a 3380 device and a target stripe count of 5 (after rounding up) on a 3390 device. An SDR of 0 results in a target stripe count of 1, and conventional volume selection occurs. A non-zero SDR which is 1 when divided by the device value results in a target stripe count of 1, but striping volume selection is used instead of conventional volume selection.

All temporary data sets with a volume count greater than one are allocated as non-striped. An SDR of 0 results in a target stripe count of 1 for both guaranteed and nonguaranteed space requests, and conventional volume selection occurs. A non-zero SDR that is 1 when divided by the device value results in a target stripe count of 1, but striping volume selection is used instead of conventional volume selection.

For non-guaranteed space, the volume must be able to satisfy the primary space requested divided by the number of stripes. For example, if primary space is 15MB, and the number of stripes is three, the volume must be able to satisfy an allocation of 5MB and not exceed this high threshold as specified in the storage group.

For guaranteed space, requests that contain specified volume serial numbers, each stripe must be able to satisfy the primary space requested (15MB in above example).

Striping and Guaranteed Space

For guaranteed space allocations, the target number of stripes is the greater of either the volume count specified or the number of specified volume serial numbers. In addition, all specified volumes must be in the same SMS storage group.

SMS assumes that the amount of space that the user wants is the target number of stripes times the primary space specified.

- If you explicitly specify volume serial numbers with guaranteed space and the target number of stripes is equal to the number of volume serial numbers you specify, SMS must allocate the primary space requested on each of these volumes. If this is not possible, the allocation fails.

- If you do not specify any volume serial numbers, then the target number of stripes is equal to the volume count. SMS tries to select the same number of volumes, but settles for less if this number is unavailable. If fewer stripes are allocated, SMS increases the allocation per volume to compensate for the fewer stripes.

- For VSAM data sets, if the number of guaranteed space volumes exceed 16 (which is the maximum number of stripes for a VSAM data component), the number of guaranteed space volumes will be reduced to 16 with the remaining volumes becoming candidates for secondary space.

- If the target number of stripes is higher than the number of volume serial numbers you specify, SMS must select all the specified volumes plus additional ones. These nonspecific volumes are not mandatory and if none are available, SMS allocates the primary quantity on each of the specific volumes.

- It is recommended that you have a sufficient number of volumes behind each controller in the storage group to reduce volume overuse. An example of a problem would be if a storage group contains eight controllers and the average stripe count is four, each controller is selected approximately 50 percent of the time. If one controller contained only two volumes, each of the two volumes is selected for approximately 25 percent of all new striped allocations.

Rules for Striped Data Sets

The maximum allocation can exceed the 64K track limit and is limited by the amount of free space below the high threshold or extent constraints mentioned above for primary space allocations.

Extending Striped Data Sets

All stripes must be able to satisfy the secondary space allocation (divided by the number of stripes) during extend processing or the allocation fails. Secondary space amount is divided by the number of stripes for both guaranteed and non-guaranteed space requests. It may be rounded up. OPEN, CLOSE, and EOV perform this.

Note that volume fragmentation information is not available to SMS at volume selection time. As a result, SMS may attempt to select a volume which appears to have enough free space but which is too fragmented to support the allocation. If this occurs, the volume becomes rejected and SMS performs striping volume reselection.

If All Else Fails....

You can specify attributes on Page 3 of the Data Class Define panel to indicate whether or not to retry new data set allocations or extends on new volumes that fail due to space constraints.

During allocation, there might not be enough space on a volume to meet the requested space. SMS volume selection can sometimes solve this problem by trying all candidate volumes before failing the allocation. You can also use the Space Constraint Relief and Reduce Space Up To (%) attributes to request that an allocation be retried if it fails due to space constraints. SMS retries

the allocation by combining any of the following:

Spreading the requested quantity over multiple volumes
Allocating a percentage of the requested quantity
Using more than 5 extents

Space Constraint Relief specifies whether or not to retry an allocation that was unsuccessful due to space constraints on the volume. Note that allocation is attempted on all candidate volumes before it is retried. This attribute applies only to system-managed data sets, and is limited to new data set allocations, and while extending the data set on new volumes.

Y specifies that SMS retry the allocation.
N specifies that SMS does not retry the allocation, so that allocation is not attempted on multiple volumes. This is the default.

Reduce Space Up to (%) specifies the amount by which you want to reduce the requested space quantity when the allocation is retried. You must specify Y for the Space Constraint Relief attribute. Valid values are 0 to 99.

If you request space constraint relief but do not specify a percentage value (either 0 or blank), SMS does not reduce the requested space quantity. This implies your application cannot tolerate a reduction in the space to be allocated, so only the 5 extent limit is relieved.

If You Use Space Constraint Relief....

- The volume count can be specified in the data class or in the JCL.
- Existing data sets may get smaller extents than requested when they extend.
- New data sets may get less primary space allocated than was requested.
- Since more extents may be used during the initial allocation, it my make less available for future extends of the data set. For example, if the primary space allocation used 10 extents when allocating a physical sequential data set, then only 6 extents are left for allocation of secondary space.

Compatibility PTFs are required for DFSMS/MVS V1R3 or below and for all DFP releases so that lower-level systems can access all 255 extents of a VSAM component created on a DFSMS/MVS V1R4 system. Note that this access is for input only. Do not assign data classes with space constraint relief for VSAM data sets if there is a possibility they might be accessed for output on a lower-level system.

The Retry Process....

The volume count canbe specified in the data class or in the JCL.
Existing data sets may get smaller extents than requested when they extend.

New data sets may get less primary space allocated than was requested.

Since more extents may be used during the initial allocation, it my make less available for future extends of the data set. For example, if the primary space allocation used 10 extents when allocating a physical sequential data set, then only 6 extents are left for allocation of secondary space.

Compatibility PTFs are required for DFSMS/MVS V1R3 or below and for all DFP releases so that lower-level systems can access all 255 extents of a VSAM component created on a DFSMS/MVS V1R4 system. Note that this access is for input only. Do not assign data classes with space constraint relief for VSAM data sets if there is a possibility they might be accessed for output on a lower-level system.

Requesting Assistance

Volume selection is a complex area, and you may not be able to determine what is going on. If you are unable to determine the cause of a specific volume selection problem, it is recommended that you try to collect some data before calling IBM support.

1. The SETSMS command turns SMS tracing on for both error and non-error conditions with a trace table size of one megabyte. It is important that the DESELECT and SELECT keywords be specified in the order shown. Be sure to include the job name to prevent the trace table from filling up with unwanted data, and possibly adversely affecting performance.

2. Once tracing is set, run the job which seems to be having volume selection problems. Running it with tracing on will cause module to module information to be captured as well as specific information on new data set creation and old data set allocation.

3. The trace should be turned off once the job completes to prevent unnecessary data from being written to the trace table.

4. Making a note of the data set name and providing it to IBM support can assist support personnel in locating the relevant trace entries.

5. The SMS trace table resides in the SMS address space. In order to process it using IPCS, a dump needs to be taken which include the private area (RGN) of the address space.

6. Once you have the dump, you can use IPCS to view it.

7. You'll need to provide the name of the data set containing the dump.

8. IPCS option 6 allows you to enter commands.

9. The VERBX command will format and display the data in the SMS trace table.

10. Printing a hardcopy of the trace table will allow it to be faxed to IBM support and may allow support to provide you with better service. Refer to the IPCS Users Guide for information on allocating a print data set with DDNAME IPCSPRNT.