



**z/OS® and DFSMS™**

***System-managed tape overview***

@business on demand software

© 2008 IBM Corporation

This education topic provides a high-level overview of the system-managed tape (SMStape) support in System z (z/OS) and ends with a discussion of the mount flow for scratch and private allocations and the components that are involved.

## Agenda

- Overview of SMStape
- SMStape components
- Scratch mount path
- Specific mount path
- Recovery action

This presentation begins with an overview of the system-managed tape (SMStape) environment. It then moves on to discuss which System z components are involved with SMStape and what role each component plays. From there, two examples are covered that will explain in detail how both a scratch and specific mount request are satisfied by the system. The presentation concludes with a short discussion on error handling and recovery action.

## SMStape overview

- The system-managed tape (SMStape) support was introduced in 1993 with the IBM 3495 Tape Library Dataserver
  - ▶ Expanded to support additional automated, virtual and manual tape libraries
- The SMStape support manages the tape volumes and the tape libraries.
  - ▶ The application and the tape management system must still manage the data
- Components involved:
  - ▶ OAM, SMS, ISMF, Catalog, AMS, O/C/EOV, Checkpoint/Restart, AOM, Device Services, DFSMSHsm™ (HSM), DFSMSrmm™ (RMM), MVS allocation, IOS/DDR, Console Services, Scheduler, JES3, plus tape ISV support.



System-managed tape (SMStape) was introduced in 1993 along with the IBM 3495 Tape Library Dataserver, IBM's first automated tape library. Over the years, additional enhancements have been made to support new automated tape libraries along with virtual and manual tape libraries.

SMStape support is responsible for managing the tape volumes and libraries in a System z (z/OS) environment. This includes tracking volume categories, determining which volumes, devices, and libraries can satisfy a request, providing information about the tape library environment, error handling and recovery actions, and many other functions. However, the host application and tape management system are still responsible for managing the data that resides on the volumes.

SMStape consists of multiple components all working together to manage the tape environment.

## SMStape overview

- Tape library is set of tape volumes and the tape drives where the volumes can be mounted.
- Exclusive relationship:
  - ▶ library-resident volume can only be mounted on the drives defined to that library.
- Tape library can be automated, virtual or manual.
- Each tape library supports multiple device types and media types.



A tape library consists of a set of tape volumes and tape drives where those volumes can be mounted. Volumes can only be mounted on tape devices defined to the library that owns the volumes. A tape library can be automated, virtual, or manual. An automated tape library uses robotics to automatically retrieve and mount a volume that is requested. A manual library requires an operator to manually mount a requested volume. A virtual tape library uses logical volumes, which reside on cache instead of physical tapes to satisfy requests. These logical volumes can then be migrated to physical tape by the library to conserve cache. If a request is made for data that resides on a migrated volume, the logical volume is written back to cache before the request is satisfied.

Each type of tape library can support multiple device types, recording formats, and media types.

## SMStape components

- **Catalog**
  - ▶ User catalog (VOLCAT) houses volume and library records.
    - Referred to as the tape configuration database (TCDB).
  - ▶ Access Method Services (IDCAMS) create, alter, delete and list support for the volume and library records (mainly for recovery-only purposes)
- **MVS Allocation**
  - ▶ Used in the allocation of job resources (for example, devices). Responsible for allocating a library device from a list of eligible devices.
- **Open/Close/End of Volume**
  - ▶ Used during data set open and volume verification processing. Provides set of installation exits that interface with the tape management systems.

SMStape support consists of these components:

Catalog is responsible for the user catalog (VOLCAT), which houses volumes and tape library records. It is commonly referred to as the tape configuration database (TCDB). Catalog also provides Access Method Services (IDCAMS), which can create, alter, delete, and list volume and library records. Access Method Services is mostly used for recovery purposes.

MVS Allocation is responsible for allocating a library device from a list of eligible devices.

Open/Close/End of Volume is used during data set open and volume verification processing to provide a set of installation exits that interface with tape management systems. These installation exits provide a means for a user to allow or disallow the use of a volume.

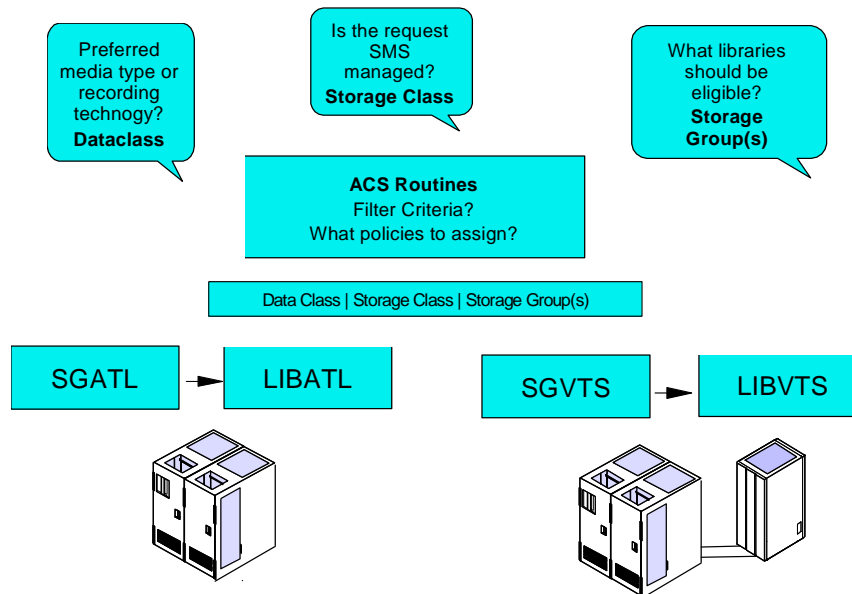
## SMStape components

### ▪ SMS

- ▶ Provides interfaces to control tape library volume and device selection.
- ▶ For new allocations:
  - ACS (automatic class selection) routines invoked during job processing
    - IF-THEN-ELSE logic used to assign policy constructs
  - Storage class ACS routine used to make request SMS-managed.
  - Storage group ACS routine used to direct the allocation to a tape libraries.
  - Data class ACS routine used for media and recording technology preferencing.

SMS provides a set of interfaces that allow a user to control tape library volume and device selection. For new allocations, the Automatic Class Selection (ACS) routines are invoked during job processing. Simple IF-THEN-ELSE logic is used to assign policy constructs to a new allocation. The three different constructs in the ACS routines used to direct the allocation are storage class, storage group, and data class. The storage class construct is used to make the request SMS-managed. The storage group construct can direct the request to specific libraries. Finally, the data class construct can be used to direct the request to a specific media type or recording technology or both. These three constructs can be specified directly in the JCL or the ACS routines can be written so that the proper construct is automatically determined after the job is submitted. In addition to storage, class, storage group and data class there is also management class. Management class is not directly used in the allocation path by the SMStape support.

## Example: ACS routines



This slide presents an example of how ACS routines are used during a new allocation request. First, the Storage Class construct is looked at to determine if this request should be SMS-managed. If the request should be SMS-managed, the Storage Group construct is then queried to determine which library should handle the request. In this example, SGATL is the storage group name assigned to an ATL library while SGVTS is the storage group name assigned to a VTS. Depending on which Storage Group is chosen, this request could be directed to either an ATL or VTS library. Finally, the Dataclass construct is queried to determine whether or not a particular media type or recording technology should be used for this request.

## SMStape components

- Device services/AOM
  - ▶ Provides interface to send I/O to library.
  - ▶ Builds and runs channel programs to the library.
  - ▶ Isolates I/O to library in a single component.
  - ▶ Maps what device pools (tape subsystems) reside in which libraries.
    - A device pool is a string of tape drives attached to a single control unit.
      - Up to 16 device per control unit image.
    - All devices in a device pool must have the same characteristics.
      - Support the same media types and recording formats.

Device Services/AOM is the component responsible for providing an interface to send I/O to a tape library. This isolates all I/O to tape libraries into a single component. In addition to building and executing channel programs to the library, Device Services/AOM also maps which device pools (or tape subsystems) reside in which tape libraries. A device pool is a string of up to 16 tape devices attached to a single control unit. All devices in a device pool must support the same media types and recording formats.



## SMStape components

### ▪ OAM

- ▶ Provides a set of library and volume-related services
- ▶ Library Automation Communication Services (**CBRXLACS**) for mount, demount, and so on.
  - communicates to library using device services
- ▶ Library Control Services (**CBRXLCS**):
  - CUA - Change use attribute (scratch/private transitions).
  - EJECT - Eject a volume.
  - QVR - Query volume residency and obtain volume information.
  - TVE - Test volume eligibility.
  - MCE - Manual cartridge entry.
  - OLN - Obtain library name.
  - And more.

OAM provides a set of library and volume related services. The Library Automation Communication Services (CBRXLACS) is primarily responsible for mounting and demounting a volume. The Library Control Services (CBRXLCS) provides an interface for users to control various aspects of the library and obtain tape library and volume related information.

## SMStape components

- OAM (continued)
  - ▶ Builds and invokes the SVC 26 – catalog interface
    - **CBRXLIB** - Creates, retrieves, updates library record.
    - **CBRXVOL** - Creates, retrieves, updates, replaces or deletes volume record.
  - ▶ Installation exits provided at key processing points:
    - **CBRUXENT** - Cartridge entry.
    - **CBRUXEJC** - Cartridge eject.
    - **CBRUXCUA** - Change use attribute processing (scratch/private transitions).
    - **CBRUXVNL** - Volume not in library (used to direct allocation into a library for shelf resident volumes).

In addition to the service interfaces, OAM also provides an interface to modify the catalog or TCDB. CBRXLIB is used to create, retrieve, and update a tape library record while CBRXVOL is used to create, retrieve, update, replace, or delete a tape volume record.

OAM also provides a set of installation exits at key processing points. CBRUXENT is used to allow or disallow the entry of tape cartridges into a tape library while CBRUXEJC is used to allow or disallow cartridges to be ejected from a tape library. CBRUXCUA is invoked to determine if a change use attribute (scratch/private transitions) is allowed for a particular volume. CBRUXVNL is used to direct allocation into a specific library for shelf resident volumes.

## SMStape components

- OAM (continued)
  - ▶ OAM address space:
    - Cartridge entry and eject processing.
    - Operator command processing.
    - Audit processing.
    - Unsolicited attention message processing (library operator messages, operational state change, device availability, category state change, and so on.).
  - ▶ Job processing (mount/demount activity) occurs outside the OAM address space.

The OAM address space must be active in order to allow cartridge entry and eject processing, operator command processing (such as library and tape volume displays), audit processing, and unsolicited attention message processing. Job processing (mount/demount activity) does not require the OAM address space to be active, however the recommendation is to always run with the OAM address space active.

## SMStape components

- RMM
  - ▶ IBM's tape management product (removable media manager):
    - manages the tape volumes and the datasets on the tape volumes
    - maintains detailed volume and dataset information in its own database
    - Tracks the volume's owner and expiration information
    - provides utilities to return volumes to scratch
    - uses the OAM tape library interfaces and exits
    - also interfaces with the Open/Close/EOV interfaces
    - tracks the movement and location of the volumes on and off-site.

RMM is a tape management product that can manage tape volumes and the datasets that reside on those volumes, and maintain detailed volume and dataset information within its own database. RMM also provides utilities to return volumes to scratch and can track the movement and location of volumes both on and off-site.

## Scratch mount (DISP=NEW) – JES2 (Part 1)

- MVS Allocation calls SMS to invoke ACS routines.
- ACS routines assign data class, storage class, management class and tape storage groups.
- Data class specifies media and recording technology preferencing.
- SMS invokes OAM with storage groups and assigned data class information.
- OAM maps storage groups to list of libraries.
- OAM invokes device services to get list of device pools.
- OAM eliminates ineligible device pools.

Now, two examples will be covered in depth to explain how both a scratch and specific mount request are satisfied by SMStape.

The first example will cover a scratch mount request in a JES2 environment where a volume is being written from load point (DISP=NEW).

The first step is MVS Allocation calling SMS to invoke the ACS routines. The ACS routines will then assign a data class, storage class, management class, and storage group to the request. These assigned constructs will determine if the request is SMS-managed, whether a particular media type or recording format is to be used, and which library will receive the request. After assigning each of the constructs, SMS then invokes OAM with the storage group and data class information. OAM then determines which libraries are associated with the assigned storage group and invokes Device Services to get a list of device pools that reside in that library. OAM then determines which devices are eligible to satisfy the request by using the media type and recording format specified in the data class construct.

## Scratch mount (Part 2)

- OAM returns list of device pools ordered by scratch volume availability (above or below threshold) to SMS.
- SMS returns list of device pools to MVS Allocation.
- MVS Allocation builds an eligible device list taking into account device type and cartridge loader status.
- MVS allocates a device.
  - ▶ For a non-deferred mount, MVS Allocation invokes OAM to mount a volume.
  - ▶ For a deferred mount, OPEN invokes OAM to mount a volume.
- With advanced policy management, the ACS routine assigned constructs are sent to the library enabling outboard policy actions (for virtual tape).

After receiving the device pools from Device Services, OAM then returns this list to SMS ordered by scratch volume availability. SMS then returns this list of device pools to MVS Allocation. MVS Allocation then builds an eligible device list taking into account both the device type and cartridge loader status. Based on this eligible device list, MVS Allocation can then allocate a device to satisfy the request. In the event of a non-deferred mount, MVS Allocation invokes OAM's CBRXLACS services to mount the volume, whereas in a deferred mount situation, OPEN/CLOSE/End of Volume will invoke OAM to mount the volume.

In the event this request is satisfied by a VTS, the ACS routine assigned constructs are then passed to the library allowing for outboard policy actions against the volume.

## Scratch mount (Part 3)

- OAM invokes Device Services to send category mount order to library.
- Library selects a scratch volume from the specified category and schedules the mount.
- MVS Allocation completes and job step starts.
- Job issues an OPEN for the tape dataset.
- OPEN invokes OAM to wait for mount to complete.
- Mount completes, library notifies Device Services, which notifies OAM, which notifies OPEN.
- OPEN reads volume label and performs label verification.

After OAM has been called to process the mount request, OAM then invokes Device Services to send the category mount order to the library. The library then selects a scratch volume from the specified scratch category and schedules the mount. MVS Allocation completes as soon as the mount is scheduled and the job step begins. The job issues an OPEN for the requested tape dataset and OPEN invokes OAM to wait for the mount to complete. When the mount completes, the library notifies Device Services, which notifies OAM, which notifies OPEN. OPEN can then read the volume label and perform label verification.

## Scratch mount (Part 4)

- If OPEN is successful, OPEN invokes SMS to move the volume to PRIVATE status.
  - ▶ Tape management system is given opportunity to approve the volume's category change.
- SMS invokes OAM to change the use attribute of the volume
- OAM invokes the change use attribute installation exit (CBRUXCUA), notifying tape management system.
- OAM invokes Devices Services to change the volume category to PRIVATE at the library.
- OAM moves the volume from Scratch -> Private in the TCDB.
- OPEN returns to caller with successful OPEN.
- Data can then be written.

If OPEN is successful, it invokes SMS to move the volume from SCRATCH to PRIVATE status. The tape management system is given an opportunity at this point in time to approve or disapprove of the volume's category change. SMS then invokes OAM to change the use attribute (CUA) of the volume using CBRXLCS. When invoked for a CUA, OAM invokes CBRUXCUA, the change use attribute installation exit, which notifies the tape management system that the volume will be changing categories. To complete the CUA request, OAM invokes Device Services to change the volume category to PRIVATE at the library. After that is complete, OAM then moves the volume from SCRATCH to PRIVATE in the TCDB. After the volume has been successfully moved to the PRIVATE category, OPEN can finally return to the caller with a successful OPEN, at which point data can be written to the volume.



## Specific mount – JES2 (Part 1)

- MVS Allocation invokes SMS to determine if volume resides in a tape library.
- SMS invokes OAM to retrieve volume record from TCDB.
- If volume is not library resident, OAM invokes the volume not in library exit (CBRUXVNL).
  - ▶ Provides an operator the opportunity to enter a shelf resident volume.
- With the volume library resident, SMS invokes OAM with a list of VOLSERS.
- OAM ensures all volumes are PRIVATE and in same storage group.

The second example details how a specific mount request is handled by SMStape and its related components.

MVS Allocation begins by invoking SMS to determine if the specified volume resides in a tape library. SMS invokes OAM to retrieve the volume record from the TCDB that can provide this information. If the volume is not library resident, OAM invokes the volume not in library exit (CBRUXVNL) to provide an operator the opportunity to enter the shelf resident volume into the library. Once the volume is determined to be library resident, SMS invokes OAM with the VOLSERS to verify that all volumes are in the PRIVATE category and belong to the same storage group.

## Specific mount (Part 2)

- OAM invokes Device Services to obtain list of devices in library.
- OAM uses volume record media and recording technology information to eliminate ineligible devices.
- OAM returns device pool list to SMS, which passes it back to MVS Allocation.
  - ▶ For read requests that can be satisfied by multiple device types, OAM returns a device pool list ordered from oldest to newest device type. This allows newer device types to stay available for requests that can only be satisfied by a newer device.
- MVS Allocation builds eligible device list and allocates a device.
  - ▶ For a non-deferred mount, MVS Allocation invokes OAM.
  - ▶ For a deferred mount, OPEN invokes OAM.

OAM then calls Device Services to obtain a list of devices in the library associated with the storage group assigned to the volume. Using media type and recording technology information, OAM eliminates ineligible devices. Once a list of eligible device pools is created, it is returned to SMS, which passes it back to MVS Allocation. Note, for a read request, OAM will order this device pool list from oldest device type to newest. This allows newer device types to stay available for requests that can only be satisfied by a newer device type. Once MVS Allocation receives the device pool list, it builds a list of eligible devices and allocates a device. As with a scratch request, in the case of a non-deferred mount, MVS Allocation invokes OAM for the mount request and in the case of a deferred mount, OPEN invokes OAM for the mount.

## Specific mount (Part 3)

- OAM invokes Device Services to send the mount order to the library.
- Library schedules mount.
- MVS Allocation completes and job step starts.
- Job issues an OPEN for the dataset.
- OPEN invokes OAM to wait for the mount to complete.

Once OAM has received the mount request, it calls Device Services to issue the mount request to the library. Once the mount has been scheduled, MVS Allocation completes processing and the job step begins. The job issues an OPEN for the dataset and OPEN in turn calls OAM to wait until the mount is complete.

## Specific mount (Part 4)

- Mount completes, library notifies Device Services, which notifies OAM, which notifies OPEN.
- OPEN reads volume label and performs label verification.
- If OPEN successful, OPEN invokes SMS to update fields in the TCDB volume record.
- SMS invokes OAM to update the last mounted date, and so on ...
- OPEN returns to caller with successful OPEN.
- Data can then be read or written.

Once the mount is complete, the library sends notification to Device Services, which notifies OAM, which notifies OPEN. OPEN then reads the volume label and performs label verification. If successful, OPEN invokes SMS to update fields in the TCDB volume record. SMS then calls OAM's CBRXVOL interface to update the volume TCDB record. Once the volume record is updated, OPEN returns to the caller with a successful notification and data can then be read or written to the volume.

## Scratch and specific mount recovery action

- If the mount cannot be scheduled or does not complete, OAM (CBRXLACS) recovery examines the error and attempts recovery.
  - ▶ OAM issues CBR4195I and CBR4196D.
    - Operator is given the opportunity to retry or cancel the mount.

```
CBR4195I LACS retry possible for job ATNCMP: 015
IEE763I NAME= CBRLACS CODE= 140169
CBR4000I LACS MOUNT permanent error for drive 0BCA.
CBR4105I No MEDIA5 scratch volumes available in library ATL10001.
IEE764I END OF CBR4195I RELATED MESSAGES
*06 CBR4196D Job ATNCMP, drive 0BCA, volser SCRTCH, error code 140169.
Reply 'R' to retry or 'C' to cancel.
```

In either scenario, if the mount request cannot be scheduled or does not complete, OAM's CBRXLACS recovery path examines the error and attempts recovery. CBR4195I and CBR4196D are issued to the console. These messages give the operator an opportunity to either retry or cancel the mount. An example of these messages is shown.

## Scratch mount recovery action

- During a scratch mount, if OPEN fails the mount, it invokes OAM (CBRXLACS) to demount volume, assign it to error category, and mounts another scratch volume (up to five additional retries).

During a scratch mount request, if OPEN fails the mount, OAM's CBRXLACS function is invoked to demount the volume, assign it to the error category, and then retry the mount with another scratch volume, up to five additional times.

## Publications

- *z/OS DFSMS Object Access Method Planning, Installation and Storage Administration Guide for Tape Libraries (SC35-0427)*

For additional information refer to the *z/OS DFSMS Object Access Method, Planning, Installation and Storage Administration Guide for Tape Libraries (SC35-0427)*.

## Feedback

### Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

[mailto:iea@us.ibm.com?subject=Feedback\\_about\\_V1R0\\_OAM\\_SMSTape.ppt](mailto:iea@us.ibm.com?subject=Feedback_about_V1R0_OAM_SMSTape.ppt)

This module is also available in PDF format at: [..V1R0\\_OAM\\_SMSTape.pdf](..V1R0_OAM_SMSTape.pdf)

You can help improve the quality of IBM Education Assistant content by providing feedback.



## Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DFSMS DFSMSshm DFSMSrmm IBM z/OS

A current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2008. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.