Program Product

OS/VS2 MVS Hierarchical
Storage Manager:
General Information

IBM

Program Product

OS/VS2 MVS Hierarchical Storage Manager: General Information

Program Number 5740-XRB
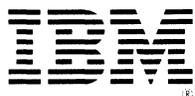
IBM

Preface

This manual is an introduction to the Hierarchical Storage Manager. It is written for data processing managers, operations managers, system programmers, applications programmers, system operators, TSO terminal users, and other data processing personnel who are involved in data space management. This publication contains general information about the Hierarchical Storage Manager, its functional characteristics, hardware considerations, operating system considerations, planning and installation, and the responsibilities that go along with this program product. Readers should have some familiarity with OS/VS2 MVS concepts and terms.

Major Divisions of This Publication

The publication is divided into the following chapters:

"Introduction," which gives an overview of the Hierarchical Storage Manager.

"Description of the Functions and Control Data Sets," which describes the functions of the Hierarchical Storage Manager and the control data sets.

"Planning and Installation," which describes the hardware, operating system, and other preinstallation considerations and the steps that are taken to install the Hierarchical Storage Manager.

"Using the Hierarchical Storage Manager," which describes the responsibilities of the system programmer, the user, and operations personnel when they use the Hierarchical Storage Manager.

This publication is prerequisite reading for all the other Hierarchical Storage Manager publications.

Related Publications

The following publications will be available when the Hierarchical Storage Manager is available.

- *OS/VS2 MVS Hierarchical Storage Manager System Programmer's Reference and Operations Guide,* SH35-0023, which describes how to install and operate the Hierarchical Storage Manager. This publication gives detailed syntax descriptions of all the Hierarchical Storage Manager commands available to the system programmer and system operator.

- *OS/VS2 MVS Hierarchical Storage Manager User's Guide,* SH35-0024, which describes the capabilities available to the user whose data sets are controlled by the Hierarchical Storage Manager. This publication also gives detailed syntax descriptions of all the commands available to the user.

- *OS/VS2 MVS Hierarchical Storage Manager Logic,* LY35-0026, which describes the program logic of the Hierarchical Storage Manager.

- *OS/VS2 MVS Hierarchical Storage Manager Messages, ,* SH35-0025, which describes the messages issued by the Hierarchical Storage Manager.

- *Introduction to the IBM 3850 Mass Storage System (MSS),* GA32-0028, which provides an overview of the Mass Storage System.

Contents

Figures

Introduction

The Hierarchical Storage Manager is a program product that manages and controls a hierarchy of storage devices having different costs, capacities, and access attributes. The Hierarchical Storage Manager program is a continuously running task under OS/VS2 MVS with JES2. It enables you to keep active data sets on fast-access storage devices, such as DASD, and to move the less active, older data sets to a lower-cost-per-byte device, such as the IBM 3850 Mass Storage System. The Hierarchical Storage Manager allows you to take advantage of the attributes of storage products such as the IBM 3330 Disk Storage Models 1 and 11, the IBM 3350 Direct Access Storage, and the IBM 3850 Mass Storage System while it manages data sets in both TSO and batch environments.

Many OS/VS2 MVS installations have experienced a large increase in the number of interactive terminal users who are creating new data sets and extending existing ones. Because of the interactive nature of these users, the data must be kept online and available. Because online storage is relatively expensive, many installations have established criteria and operating procedures that maintain a balance between the user's needs for storage and storage costs. For example, in order to free up storage space, user data sets that have not been accessed for a period of time can be copied onto magnetic tape so that the space they occupy can be made available for more current demands. The off-loaded data stored on reels of tape is usually kept in a tape library until it is requested. This process requires human intervention that can cause problems with data integrity, performance, operation, and cost. This process also causes problems with installation physical space.

Not only are the terminal workload and related online storage requirements increasing, but many batch jobs are relying more on permanently mounted volumes to avoid processing delays. The demand for online space is increasing because of this. Therefore, many large installations have space management personnel who are responsible for allocating online space to users and for establishing procedures to reclaim it.

The Hierarchical Storage Manager can automatically keep the active data sets on DASD and move the less active data sets to another device, such as the Mass Storage System, in both TSO and batch environments. This increases effective usage of DASD and reduces the operational problems of archiving data sets on tape and it is accomplished with little or no impact on the user. When data that has been moved is referenced, it is then automatically brought back to user-accessible storage.

The Hierarchical Storage Manager can exploit the cost relationship between DASD and the Mass Storage System by keeping all data sets under its control and making them available to the terminal or batch user.

In addition, the Hierarchical Storage Manager provides automatic backup of those data sets that have been changed since the last backup was made. The ability to provide backup support at the changed data set level, instead of at the volume level for all data sets, offers many advantages.

The user, the system programmer, and the Hierarchical Storage Manager authorized data base administrator have the ability to override the automatic functions of the Hierarchical Storage Manager with commands.

The following definitions explain terms used throughout this book in describing the Hierarchical Storage Manager.

- Primary volumes–volumes containing data sets directly accessible to the user and under control of the Hierarchical Storage Manager. These volumes are normally permanently resident on DASD. They can reside on 3330 Model 1, 3330 Model 11, or 3350 devices, or they can be mounted volumes in the Mass Storage System. Recommended use is that they be on 3330 Model 1, 3330 Model 11, or 3350 devices, or mounted, bound volumes in the Mass Storage System.

- Level 1 volumes–volumes to which the Hierarchical Storage Manager moves data sets from primary volumes. These volumes can reside on 3330 Model 1, 3330 Model 11, or 3350 devices, or in the Mass Storage System. Recommended use is that these volumes be mounted, mass storage volumes.

- Level 2 volumes–volumes to which the Hierarchical Storage Manager moves data sets from level 1 or primary volumes. This level enables the user to group data sets by high-level qualifier. These volumes can reside on 3330 Model 1, 3330 Model 11, or 3350 devices, or in the Mass Storage System. Recommended use is that they be demounted, mass storage volumes.

- Backup volumes–volumes to which the Hierarchical Storage Manager copies data sets. These volumes provide recovery and multiple data set version capability. These volumes can reside on 3330 Model 1, 3330 Model 11, or 3350 devices, or in the Mass Storage System. Recommended use is that they be demounted, mass storage volumes.

- Spill volumes–volumes to which the Hierarchical Storage Manager moves all but the most current valid backup copies of data sets, whenever additional space is needed on a backup volume. These volumes can reside on 3330 Model 1, 3330 Model 11, or 3350 devices, or in the Mass Storage System. Recommended use is that they be demounted, mass storage volumes.

The Hierarchical Storage Manager has four major functions that provide space management, backup, and recovery capability. These four functions are (1) data migration, which moves less active data sets off primary volumes, (2) recall, which moves migrated data sets back to primary volumes, (3) backup, which copies data sets to backup volumes, and (4) recovery, which recovers data sets from the backup volumes. These functions occur automatically or can be started with a command, except for recovery, which can only be started with a command. The automatic processing is controlled by several parameters that are specified by the installation.

The Hierarchical Storage Manager maintains control data sets that contain information about the processing done by its major functions. These control data sets are updated as activity occurs in the Hierarchical Storage Manager. The Hierarchical Storage Manager provides the ability to back up and recover these control data sets. It also provides the ability to produce reports about the processing done by the Hierarchical Storage Manager's major functions.

Description of the Functions and Control Data Sets

When the data migration, recall, backup, and recovery functions are performed, information about data sets and volumes that have been processed with these functions is recorded in the Hierarchical Storage Manager control data sets.

This chapter describes the major functions of the Hierarchical Storage Manager and the control data sets that contain information about the processing of the major functions.

Major Functions

With the Hierarchical Storage Manager, you can enhance and control online DASD space management with the following major functions:

* Migration—which moves data sets either from primary volumes to level 1 or level 2 volumes or from level 1 volumes to level 2 volumes

* Recall—which moves migrated data sets back to primary volumes

* Backup—which copies data sets to backup volumes

* Recovery—which restores copies of previous versions of data sets from backup volumes

Except for the recovery function, these major functions can be set up to occur automatically without user intervention. The recovery function can only be started explicitly with a command. The other major Hierarchical Storage Manager functions can also be started with a command. This can be done by the user, by the system programmer, or by the Hierarchical Storage Manager authorized data base administrator.

Migration

Migration is the process of moving data sets from primary volumes and placing them on level 1 or level 2 volumes or from level 1 to level 2 volumes. The purpose of migration is to keep primary and level 1 volumes as free of less active data sets as needed and to provide space for future allocations. Migration occurs automatically when invoked by conditions specified by the installation. Migration can also be invoked by a command.

There are two types of automatic migration: general and interval. *General migration* occurs at a specified time of day selected by the installation.

Interval migration occurs at a specified regular time interval during the day. At that time, interval migration checks the space on each primary and level 1 volume controlled by the Hierarchical Storage Manager.

There are four migration parameters that are specified for use by the two types of automatic migration: time of day, time interval, thresholds of occupancy, and age limit of data sets.

The *time of day* indicates to the Hierarchical Storage Manager what has been determined to be the best time to do general migration on the volumes under Hierarchical Storage Manager control.

The *time interval* is how often interval migration checks the space on primary volumes and level 1 volumes.

The *thresholds of occupancy* are the limits of space to be occupied on a volume under Hierarchical Storage Manager control. Each primary, level 1, and backup volume under control of the Hierarchical Storage Manager can have its own threshold values.

The *age limit* of data sets is specified in days and indicates to the Hierarchical Storage Manager the number of days that a data set must not have been accessed in order for it to be eligible for migration, or the data sets that are to migrate when no thresholds are specified for a volume. The age of a data set is the number of days since the data set was last referenced.

If a primary volume has thresholds of occupancy, both general and interval migration use the thresholds when migrating data sets from the volume. General migration processes the volume if the amount of occupied space on the volume is greater than the low threshold. Interval migration processes the volume only if the high threshold is met or exceeded.

Consider a primary volume whose high threshold of occupancy is set at 90 percent, and whose low threshold of occupancy is set at 80 percent. General migration processes the volume if the volume is greater than 80 percent full. Interval migration processes the volume if the volume is equal to or greater than 90 percent full. When the primary volume is processed, data sets migrate, in the order of oldest data sets first, until the low threshold of occupancy is reached. In this example, automatic migration stops when 20 percent or more of the volume is free.

If a volume has no thresholds of occupancy, general migration uses the age limit of data sets when moving data sets from the volume. For example, if the age limit is seven days, all data sets that have not been referenced for seven or more days migrate. Interval migration does not process a volume that has no thresholds.

By periodically checking the space status of primary volumes, the Hierarchical Storage Manager can manage these volumes as directed by the installation. By setting the thresholds of occupancy and the age limit for data sets, the installation can help provide the amount of storage on primary volumes needed for the users to create new data sets and extend active data sets. The Hierarchical Storage Manager manages the migration of data sets from level 1 to level 2 volumes. This also relieves the users from having to do space management.

The user can explicitly specify migration for a data set and the level of migration storage to be used. Migration invoked by a command gives the user the ability to do additional space management, if desired. It also gives the user the ability to cause data migration from volumes that are not under Hierarchical Storage Manager control.

Another aspect of Hierarchical Storage Manager migration involves small data sets, those that contain a specified number of tracks or less of data. Small sequential data sets can be processed in a manner referred to as small data set packing. *Small data set packing* causes data sets of a specified number of tracks or less of data to migrate to a VSAM data set on a level 1 volume. For example, if the specified number of tracks is two, any data sets with two or less tracks of data would migrate, as records, to a VSAM data set. The number of tracks allocated to the data set that migrates makes no difference. If 10 tracks were allocated to a data set but only two are used, the data set migrates to the VSAM data set as records. By writing the small data sets to the VSAM data set as records, better usage of space on the migration volume is achieved by using the full track. If small data set packing is used, a VSAM data set should be defined for each level 1 volume.

An advantage of the Hierarchical Storage Manager is that it only moves data, not the entire allocated space, when data sets migrate or are recalled. By only moving data, not allocated space, unused data space is released. Partitioned data sets are compressed and user information in partitioned data set directories is maintained.

The operator and Hierarchical Storage Manager authorized data base administrator can explicitly cause migration of data sets on a volume basis, and the user can explicitly cause migration of individual data sets. Together, automatic and explicit migration provide additional space management capability to both the installation and the user.

Data Sets That Do Not Migrate

There are some data set types that do not migrate with the Hierarchical Storage Manager, either automatically or explicitly. These data sets are:

- VSAM data sets or catalogs.

- ISAM data sets.

- Unmovable data sets.

- Partitioned data sets with more than one NOTE list or with more than three user TTRs for each member.

- Multivolume data sets.

- Split-cylinder data sets.

- User-labeled data sets.

- Uncataloged data sets.

- Data sets not accessible through the standard catalog search.

- Data sets whose names begin with SYS and whose fifth character is a period, or whose names are SYSCTLG.

- Data sets with no extents.

- Data sets retained from migration by the SETMIG command.

- Data sets specified in the installation exit module ARCMIMCX as data sets that can't migrate. This situation applies to volume-oriented migration only. Specific data sets can explicitly migrate even if they have an entry in the ARCMIMCX module.

Recall is the process of bringing back to primary volumes a data set that has migrated. Data sets are recalled from both level 1 and level 2 volumes, either automatically or explicitly.

In order for automatic migration to be an effective and useful service, automatic recall is possible without placing any demands on the user. Migrated data sets are recalled automatically when a job or a TSO user references the data sets. The recalling of data sets is normally automatic without any required action from the user. However, the user can explicitly recall a migrated data set with a command. The user doesn't have to know what level of migration a data set is on in order to recall it.

If there is not enough space on the primary volumes to recall a data set, the recall must be reinitiated after primary space is made available.

Through the use of a command, data sets can migrate from a volume not controlled by the Hierarchical Storage Manager. Unless otherwise specified, the recall of these data sets will be to a primary volume.

Figure 1 shows an overview of the Hierarchical Storage Manager migration and recall functions.
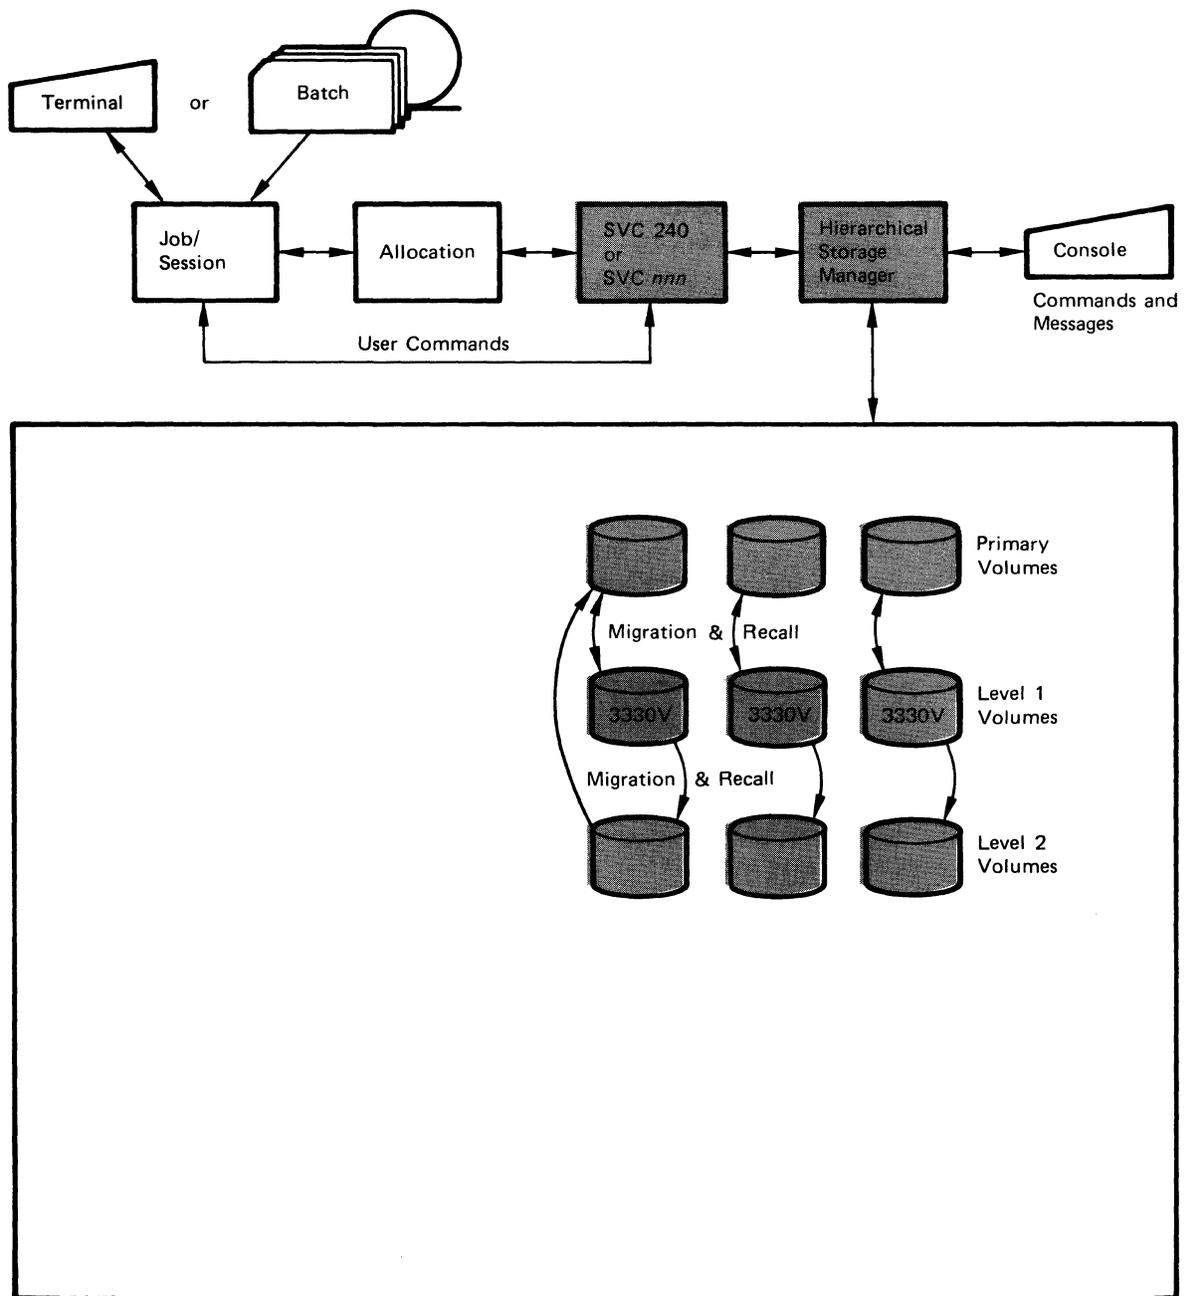
Figure 1. Migration and Recall

Backup

Backup is the process of copying data sets to backup volumes. If something happens to either the volume or a data set on the volume, the backup volumes can be used to recover the volume or data set to a specific level. Backup occurs either automatically or explicitly.

Those volumes to be controlled with automatic backup must be identified to the Hierarchical Storage Manager. The cycle, the frequency of backup, and the number of backup versions of a data set to be maintained can be specified. The *cycle* defines how often the backup function runs. The *frequency* is the number of days that must elapse since the last backup copy was made before a changed data set becomes eligible for backup. These parameters can be set up by the installation for all the users of the backup function.

An advantage of the Hierarchical Storage Manager's automatic backup function is that it makes backup copies of only those data sets that have been changed since the last backup was taken. A *changed data set* is a data set that has been opened for other than read-only access. Backing up changed data sets is called *incremental backup*. Incremental backup differs from backup utilities that copy all data on the volume on a periodic basis, whether or not the data has been changed.

For example, if your installation does daily backup using IEHDASDR and you have 25 volumes, 25 or more tapes are required each day to contain the backup copies. If you want two versions, 50 or more tapes are required for each backup run.

Initially, the Hierarchical Storage Manager will make a backup copy of each data set on all the volumes it controls with automatic backup. From this point on, automatic backup makes copies of only those data sets that have changed since the date of last backup.

If data sets occupying 25 percent of your data space change, only a size equal to 25 percent of your data space is required on the backup volumes.

If the installation specifies that two versions are to be kept, the Hierarchical Storage Manager maintains two versions of the data set on backup volumes. When space on backup volumes is insufficient, all but the latest version of each data set is moved to a spill volume.

In a Mass Storage System environment, the installation can define a sufficient number of backup volumes so that spill volumes would not be required.

The automatic backup function only processes primary and level 1 volumes. If a data set was not backed up before it migrated from a primary volume, it is still a candidate for backup from a level 1 volume.

Once the installation has designated the volumes to be controlled by the backup function, specified the cycle, specified the frequency of backup, specified the time of day, and specified the number of versions of backup to be maintained, the backup function occurs automatically.

There is no need for the user to be involved unless either a different frequency of backup or number of versions maintained is needed for individual data sets. The user can specify that individual data sets have a frequency of backup and a number of versions maintained that is different from the installation defined parameters. The user can explicitly have data sets backed up even when they are not controlled by the Hierarchical Storage Manager backup function. These capabilities give the user flexibility in data recovery in addition to the functions provided by automatic backup.

Data Sets That Are Not Backed Up

There are some data set types that are not backed up by the Hierarchical Storage Manager, either automatically or explicitly. These data sets are:

- VSAM data sets

- ISAM data sets

- Unmovable data sets with more than one extent

- Partitioned data sets with more than one NOTE list or with more than three user TTRs per member

- Multivolume data sets

- Split-cylinder data sets

- User-labeled data sets

Although data sets must be cataloged and accessible through the standard catalog search to migrate, data sets do not have to be cataloged to be backed up by the backup function.

Recovery

Recovery is the process of bringing back to a primary volume or a volume not controlled by the Hierarchical Storage Manager a backup copy or version of a data set or all the data sets that were on primary volumes or volumes not controlled by the Hierarchical Storage Manager. Recovery is done to recover data sets that have been damaged or lost, or to access an earlier version without deleting the current version.

Recovery must be initiated explicitly with a command. The user can request recovery of the latest copy or a specific version of an individual data set; the operator or Hierarchical Storage Manager authorized data base administrator can request recovery of individual data sets or all data sets on a volume.

When the need arises to recover an individual data set, the recovery function gets either the latest backup version of the data set or a specific backup version by version number or date of backup. When the need arises to recover all the data sets on a volume, the recovery function gets the latest versions of those data sets from Hierarchical Storage Manager backup volumes.

The recovery function can be supplemented with the latest full dump of the volume, and recovery can be from the date of the last dump.

A data set can be recovered from the Hierarchical Storage Manager backup volume and renamed so that the old version and the current version are both online and accessible.

Figure 2 shows an overview of the Hierarchical Storage Manager backup and recovery functions.
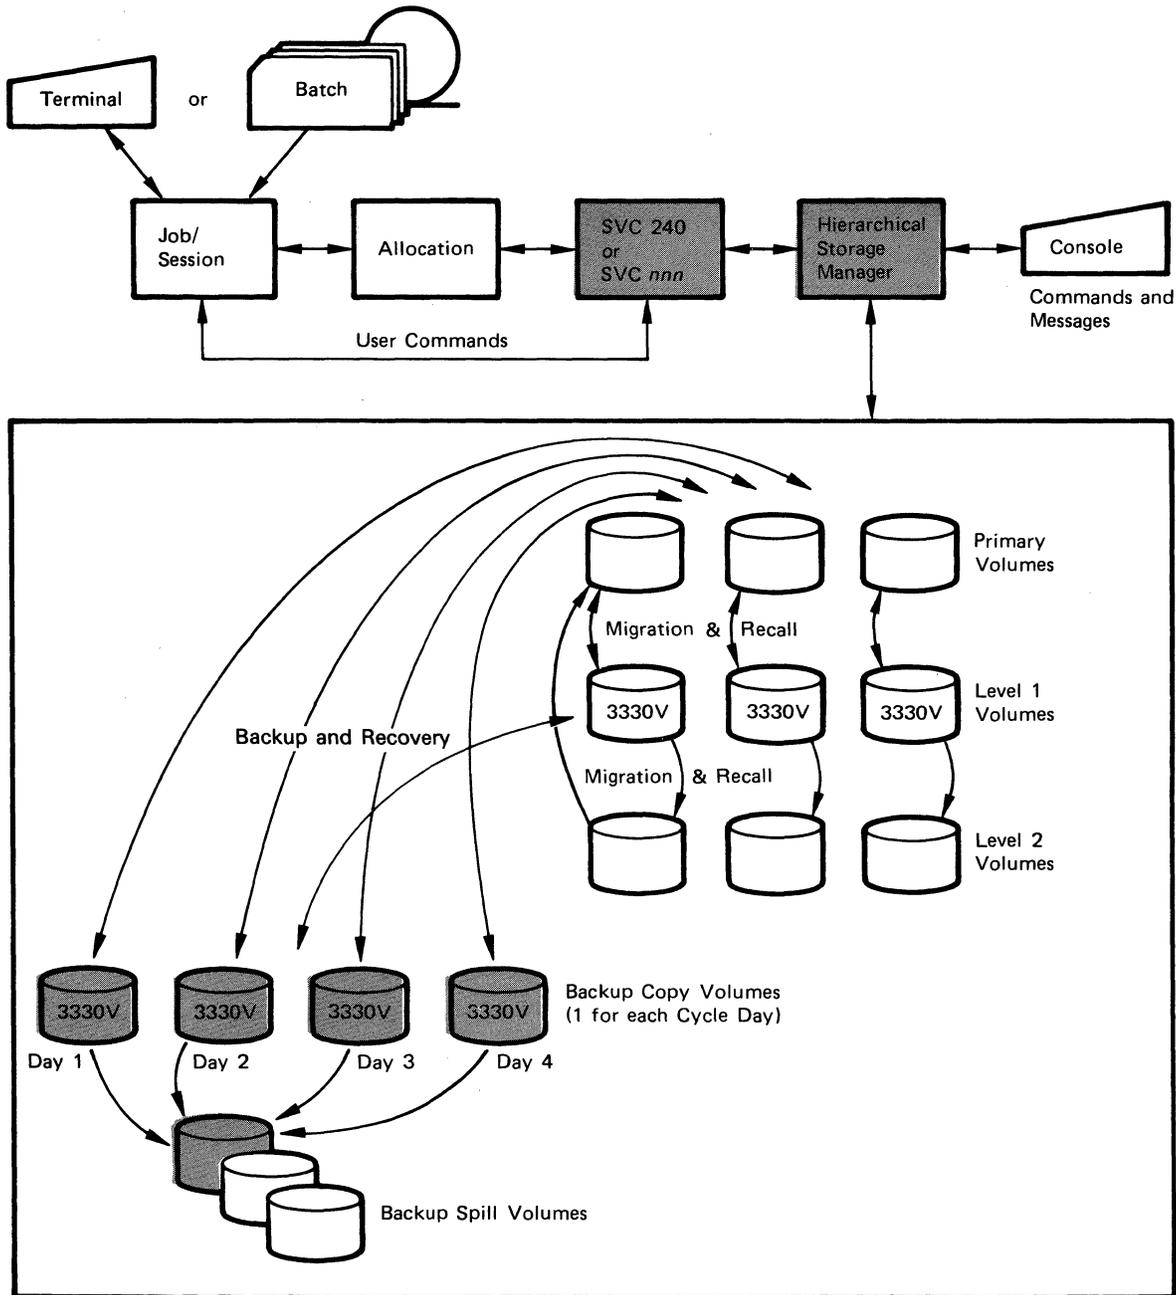

Figure 2. Backup and Recovery

Control Data Sets

All data sets under migration control of the Hierarchical Storage Manager have to be cataloged in a catalog accessible through the standard catalog search. This is essential to the Hierarchical Storage Manager migration processing. When a data set migrates or is recalled, an indicator of the migration or recall is set in the system catalog.

Records of the migration and backup activity are kept in two control data sets: the migration control data set and the backup control data set. A third data set, the journal data set, contains information necessary to recover the migration control data set and the backup control data set if either control data set is lost or damaged. The control data sets can be recovered by combining entries in the journal data set with the checkpoint copy of the affected data set. Also, transactions made in the Hierarchical Storage Manager are recorded in a pair of data sets called the log.

The rest of this chapter describes how the control data sets and the system catalog are used by the Hierarchical Storage Manager and how the control data sets of the Hierarchical Storage Manager are recovered. It contains the following topics:

- System catalog
- Migration control data set
- Backup control data set
- Journal data set
- Log
- Control data set recovery

System Catalog

The system catalog is updated when a data set migrates. The volume serial number of the volume where the data set resides is changed to MIGRAT in the catalog, which indicates that the data set has migrated. When the data set is recalled, the Hierarchical Storage Manager changes the catalog entry to reflect the data set's new location.

Migration Control Data Set

The *migration control data set* is a VSAM key-sequenced data set that is updated during migration and recall. It contains information about migrated data sets and volumes under control of the Hierarchical Storage Manager. Additionally, there is statistical and control information about migration.

The migration control data set can optionally have a backup copy. As changes are made to the migration control data set, information necessary for recovery is written in the journal data set.

Backup Control Data Set

The *backup control data set* is a VSAM key-sequenced data set that is updated during backup and recovery. It contains information about backup versions of data sets, backup volumes, and volumes under control of the backup function of the Hierarchical Storage Manager. Additionally, there is control information about backup.

The backup control data set can optionally have a backup copy. As changes are made to the backup control data set, information necessary for recovery is written in the journal data set.

Journal Data Set

The Hierarchical Storage Manager *journal data set* is a non-VSAM sequential data set. It contains information necessary for recovery of the migration control data set and the backup control data set.

Log

The Hierarchical Storage Manager *log* is a pair of non-VSAM sequential data sets. It contains records of transactions occurring in the Hierarchical Storage Manager. Additionally, some statistical information that is written in the migration control data set and the backup control data set is also written in the log.

The log consists of an X and a Y data set. When one becomes full, a swap is made to the other, and the operator is notified to dump the full data set. The full data set can be printed.

Figure 3 shows the migration control data set, the backup control data set, the journal data set, the log, and how they can reside together on volumes.
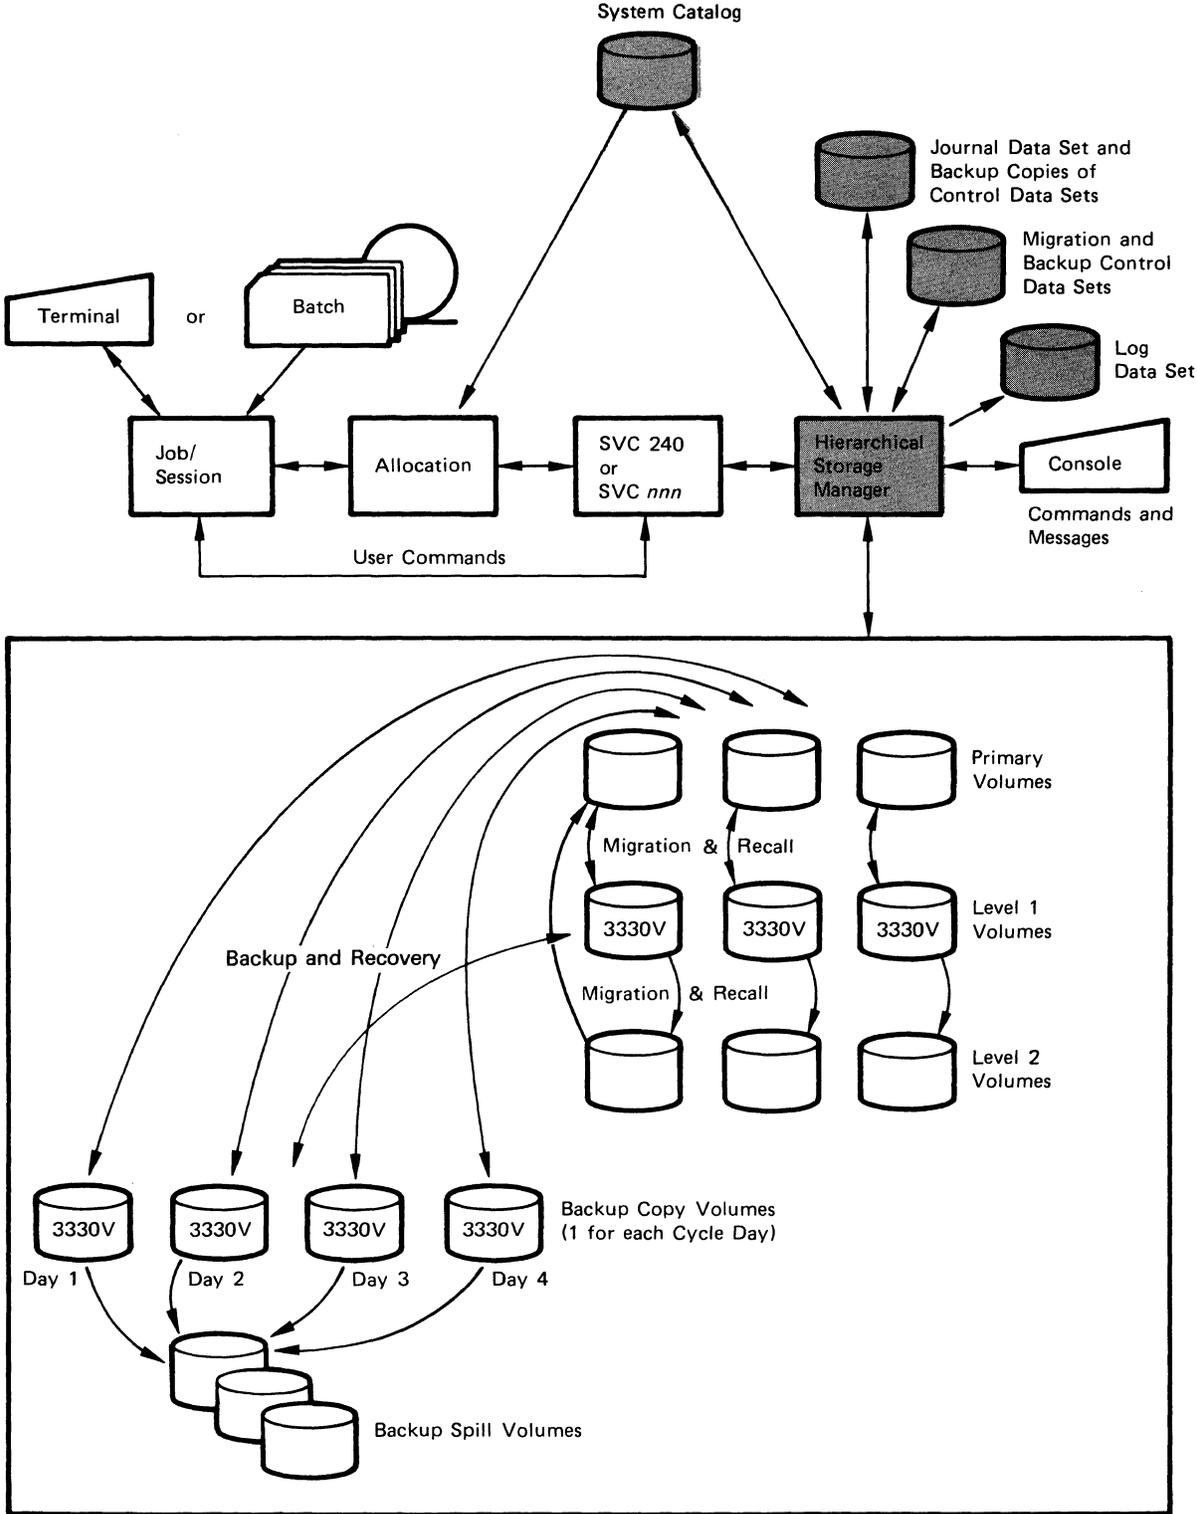
Figure 3. Control Data Sets

Control Data Set Recovery

To recover a control data set, combine the entries in the journal data set for a lost or damaged control data set with the checkpoint copy of the affected data set. Commands are provided to accomplish this process. Normal processing can be resumed after the control data sets are recovered.

It is very important that the migration control data set and backup control data set reside on a different volume from the volume where the journal data set and backup copies of the migration control data set and backup control data set reside. By doing this, the ability to recover the control data sets is further ensured.

Planning and Installation

This chapter describes the planning that should be done before installing the Hierarchical Storage Manager. It describes the hardware considerations, the operating system considerations, and the other preinstallation planning considerations. This chapter also describes considerations for installing the Hierarchical Storage Manager.

Hardware Considerations

The hardware supported by the Hierarchical Storage Manager is the 3330 Model 1, the 3330 Model 11, the 3350, and the Mass Storage System.

The Mass Storage System is the recommended device to contain migration and backup volumes. The following are some of the reasons for using the Mass Storage System.

- The Mass Storage System provides low cost storage that can be economically expanded as storage requirements grow.

- Migrated and backed up data can be accessed in the Mass Storage System without having to retrieve and mount disk packs or tape reels.

- Backed up data can be kept for a long period of time. Most data can be recovered without operator intervention. No backup DASD packs or tape reels are required. No spill volumes are required if enough mass storage volumes are defined as backup volumes. Many versions of a data set can be kept easily.

- Migrated data sets kept in the Mass Storage System may enable the installation to reduce the amount of DASD required, while still providing availability to data. Most data sets can be recalled without operator intervention. Level 1 to level 2 migration can be reduced by having more volumes for level 1.

- When using mass storage volumes as your level 1 volumes, you could eliminate the requirement for level 2 volumes by defining enough level 1 volumes so that migration from level 1 to level 2 volumes need not occur.

- Backup versions of the migration control, backup control, and journal data sets can be kept in the Mass Storage System. The log can also be kept in the Mass Storage System.

- Primary volumes can reside in the Mass Storage System. If so, they should be mounted with the bind option.

- When you use Mass Storage System mounted volumes for primary volumes, level 1 volumes, or the other volumes controlled by the Hierarchical Storage Manager, and you include entries for them in member VATLSTxx in SYS1.PARMLIB, the volumes are mounted without operator request or intervention.

For more information about the Mass Storage System, refer to the publication *OS/VS Mass Storage System (MSS) Planning Guide*.

When the installation does not have a Mass Storage System, it is important that your level 1 volumes are mounted. Otherwise, the system operator is required to mount the volumes as they are needed.

Operating System Considerations

The Hierarchical Storage Manager runs on the OS/VS2 MVS operating system, with JES2. JES3 is not supported. The Hierarchical Storage Manager maintains Resource Access Control Facility (RACF) and password protection.

If the system configuration is multihost, the Hierarchical Storage Manager can be active on more than one of the hosts, at the same time sharing the migration control, backup control, journal, and user data sets. All hosts on which the Hierarchical Storage Manager is active must have the OS/VS2 MVS operating system, and all volumes managed by the Hierarchical Storage Manager must be shared.

Storage Requirements

The Hierarchical Storage Manager will require the following estimated amounts of storage space.

- Fixed Storage: 5K bytes (5 x 10^3)

- Dynamic Storage: 350K bytes (350 x 10^3)

- Auxiliary Storage: 500K bytes (500 x 10^3)

Other Preinstallation Planning Considerations

One of the most important activities that precedes the installation of the Hierarchical Storage Manager is preinstallation planning.

The installation should examine its data set naming convention to ensure maximum benefit from the Hierarchical Storage Manager. The process of getting as many of the installation's data sets as possible under control of the Hierarchical Storage Manager may be easier with a data set naming convention.

The following items should be considered when the installation examines the data set naming convention:

- The Hierarchical Storage Manager allows the installation to define private pools of volumes to which migrated data sets are recalled. When defining a pool, the installation specifies the high-level qualifier for data sets that are recalled to the pool's volumes. Any data set with the specified high-level qualifier is recalled to one of the volumes in the pool.

- The data sets on level 2 volumes are grouped according to the highest level qualifier of the fully qualified data set name. The installation can define a specific level 2 volume to contain data sets with a unique highest-level qualifier.

Because of these two points, setting up a naming convention that separates those data sets that are not under control of the Hierarchical Storage Manager from those data sets that are under its control can be advantageous. Setting up a naming convention that separates those data sets for one application from those data sets for another application can also be advantageous.

Along with defining a data set naming convention, the installation may want to assign the volumes to groups of units. This may require a system generation (SYSGEN) to define the name given to a group of devices and the unit addresses assigned to the group. Figure 4 shows how a data set naming convention can be used with the Hierarchical Storage Manager.

Hierarchical Storage Manager
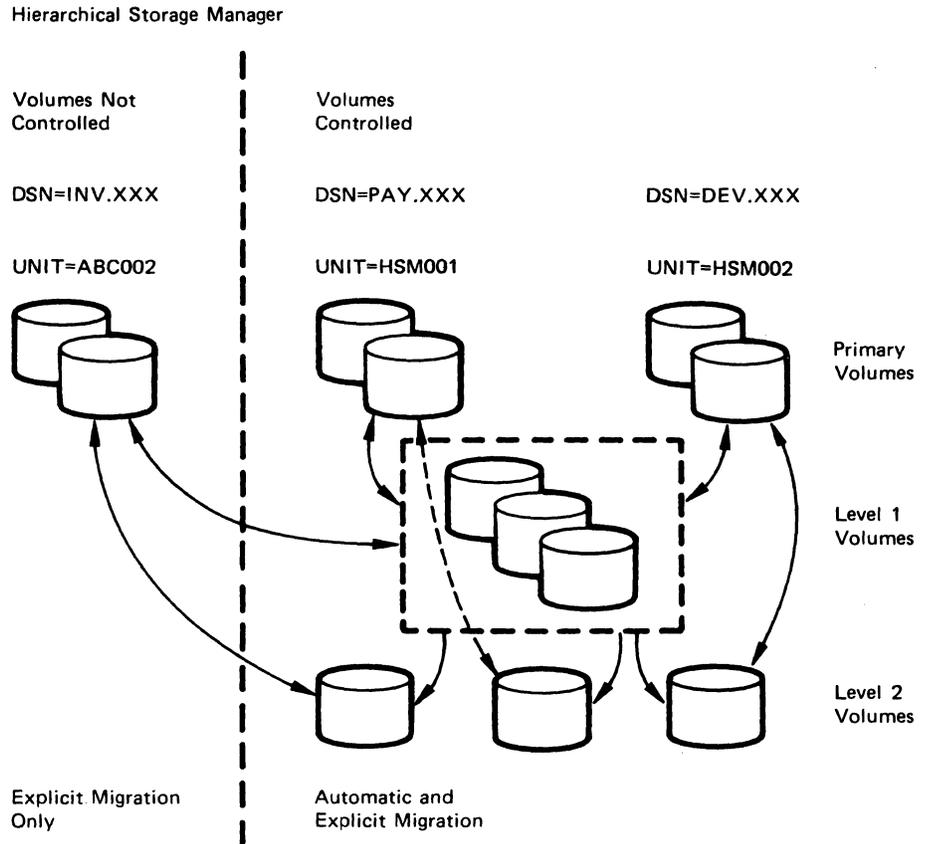
Figure 4. Migration Using A Data Set Naming Convention

A SYSGEN will be necessary to define an SVC for the Hierarchical Storage Manager, if there are no user SVCs available.

The installation should use SVC 240 for the Hierarchical Storage Manager. This SVC (a type-2, unauthorized, no-locks-held SVC) is defined in the Hierarchical Storage Manager code. If SVC 240 is already being used, the installation will have to either define an SVC or redefine an existing SVC.

The system programmer must plan for those volumes that will be controlled with automatic migration and backup. Getting volumes under Hierarchical Storage Manager control can be a volume-by-volume process by which the installation adds one volume at a time to the Hierarchical Storage Manager control. The system programmer must also plan for volumes to contain the migrated data sets and volumes to hold the backup copies. The parameters for automatic migration and backup must be planned. Careful planning and identification of the installation requirements for automatic migration and backup should be done before installation of the Hierarchical Storage Manager.

Installation Considerations

This section lists several steps that are necessary to install the Hierarchical Storage Manager. Depending on the installation, some of the steps may not have to be performed.

The following are the installation considerations:

- Install the required supporting system control program (SCP) code available in selectable unit (SU) 5752-860.

- SYSGEN SVC 240 as the Hierarchical Storage Manager SVC, or if necessary, create an SVC and change the Hierarchical Storage Manager code to reflect the new SVC number.

- SYSGEN the new unit names, if desired.

- Insert the Hierarchical Storage Manager modules into the system.

- Define the required Hierarchical Storage Manager migration control data set and backup control data set.

- Allocate the journal data set, the log data sets, and the backup copies of the control data sets.

- Define the small-data-set-packing data sets on level 1 volumes.

- Create entries for desired online volumes in member VATLST xx in SYS1.PARMLIB or set up procedures for operator mounting or demounting of these volumes.

- Set up a start-up procedure for the Hierarchical Storage Manager, which may be automated so that it occurs with system initialization.

- Set up a procedure for dumping and printing the log data sets.

- Set up a SYS1.PARMLIB member for the Hierarchical Storage Manager parameters.

For more information on defining the Hierarchical Storage Manager data sets, the start-up procedure, a procedure for the printing the log, and how to verify the installation of the Hierarchical Storage Manager, see *OS/VS2 MVS Hierarchical Storage Manager System Programmer's Reference and Operations Guide.*

Gradual Conversion of Data Sets Under Hierarchical Storage Manager Control

To ease the impact of the Hierarchical Storage Manager on the installation, the Hierarchical Storage Manager has a debug mode of operation which prevents data set movement from occurring, but which produces messages stating what would have happened if the Hierarchical Storage Manager were in normal mode.

To use the debug mode of operation, decide which data sets and which volumes that you want the Hierarchical Storage Manager to control. Then run the Hierarchical Storage Manager in the debug mode of operation to prevent data set movement from occurring. When you are satisfied that your data sets and volumes are being controlled as you would like them to be, you can run the Hierarchical Storage Manager in the normal mode. You can use this gradual conversion procedure whenever you wish to put additional data sets or volumes under Hierarchical Storage Manager control.

Using the Hierarchical Storage Manager

This chapter describes the responsibilities of the system programmer and the capabilities available to the user and the operations personnel.

System Programmer

As described in the chapter "Planning and Installation," the system programmer has several responsibilities. Basically, the system programmer's responsibilities consist of the following tasks:

- Planning for the Hierarchical Storage Manager

- Installing the Hierarchical Storage Manager

- Monitoring migration and backup and changing the migration and backup parameters as required

The system programmer has to plan for the installation of the Hierarchical Storage Manager. This planning should include examination of the installation's data set naming convention. The planning also includes determining if a SYSGEN is required, and if it is, what has to be done.

The system programmer has to install the required supporting OS/VS2 MVS SCP code necessary for use of the Hierarchical Storage Manager. In addition, the system programmer has to make the necessary SYS1.PARMLIB member entries; insert the Hierarchical Storage Manager modules into the system using the linkage editor; create the migration control, backup control, journal, and log data sets; and set up the operating procedures.

The system programmer has to define those volumes that are to be controlled by automatic migration and backup. He has to set up the installation parameters for automatic migration and backup. The migration parameters are: the times for general and interval migration, the thresholds of occupancy, and the age limit of data sets. The backup parameters are: the time of day for automatic backup, the backup cycle, the backup frequency, and the number of versions to be maintained.

If the installation has a lot of small data sets, the system programmer may want to use the small-data-set-packing facility. This will require that he define a VSAM data set on each level 1 volume and that he specify the maximum size of data sets to be migrated and recalled with this facility.

The system programmer has to authorize the Hierarchical Storage Manager data base administrators. He also has to make sure that all data sets that are to be under automatic migration control are cataloged and accessible through the standard catalog search.

The system programmer will have to monitor the migration and backup functions to ensure that the desired space management is occurring. For example, he should make sure that recently migrated data sets are not being recalled immediately.

The User

As stated throughout this book, the user is not required to issue any Hierarchical Storage Manager commands or even know that his data is being controlled with automatic migration and backup.

However, if the user wants to get involved with the data management provided by the Hierarchical Storage Manager, using commands he can:

- Change some of the backup parameters for individual data sets

- Cause a data set to migrate, to be recalled, to be backed up, or to be recovered

- Query the Hierarchical Storage Manager to determine the status of the requests that have been submitted

- List data set information from the migration control data set that indicates what data sets have been migrated and what level of migration storage they are on

- List data set information from the backup control data set that indicates what data sets have been backed up and where the backup copies are stored

- Delete a migrated data set without recalling the data set to a primary volume

Operations Personnel

The Hierarchical Storage Manager can be set up so it is started automatically when the OS/VS2 MVS system is initialized. If this is not done, the operator must enter a command to start the Hierarchical Storage Manager. Also, if for some reason the Hierarchical Storage Manager must be restarted, the operator can restart it with a start command.

If desired, during the start-up of the Hierarchical Storage Manager, several informational and statistical messages are issued. These messages advise the operator of such things as the current parameters, like the time of day for automatic migration, and the number of starts of the Hierarchical Storage Manager during the day. Messages with space information about the volumes under control of the Hierarchical Storage Manager can be issued at start-up and during interval migration when the thresholds are checked. Based on the information given in the messages, the operator may want to take some action.

There are several commands available to the operator to do certain functions. The operator can hold the migration, recall, or backup and recovery functions; write messages to the log; modify certain operating parameters; start a function that was held; and switch the X and Y log data sets.

The operator can query the status of the Hierarchical Storage Manager at any time. He can ask for information on the parameters, for statistics about operation, and for the status of pending requests.

The operator can cause data sets to migrate either on a data set or volume basis if the need arises. Also, the operator can change the Hierarchical Storage Manager parameters, such as the times for automatic migration and automatic backup. However, these activities can be covered by an installation operating procedure, or they can be done under the direction of the system programmer.

Index

The entries in this index are shown as they appear in the text of the book, that is, capitalized words are capitalized in text, lowercase words are lowercase in text, etc.

OS/VS2 MVS Hierarchical Storage Manager:
General Information
GH35-0007-1

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. This form may be used to communicate your views about this publication. They will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

IBM shall have the nonexclusive right, in its discretion, to use and distribute all submitted information, in any form, for any and all purposes, without obligation of any kind to the submitter. Your interest is appreciated.

Note: Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or the IBM branch office serving your locality.

Comments:

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Fold and tape   Please Do Not Staple   Fold and tape

First Class
Permit No. 568
Boulder,
Colorado

Business Reply Mail
No postage stamp necessary if mailed in the U.S.A.

Postage will be paid by:

International Business Machines Corporation
Information Planning and Editing, Dept. 21E
P. O. Box 1900
Boulder, Colorado 80302
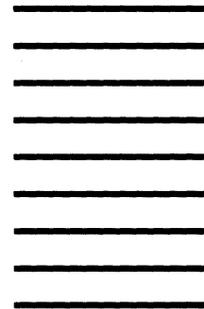
Fold and tape   Please Do Not Staple   Fold and tape

IBM ®

International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601

GH35-0007-1

IBM