IBM

MVS/ESA
Storage Administration Reference

Version 3 Release 1

IBM

MVS/ESA
Storage Administration Reference

Version 3 Release 1

Third Edition (June 1989)

This edition replaces and makes obsolete the previous edition, SC26-4514-1.

This edition applies to Version 3 Release 1 of MVS/DFP™, Program Number 5665-XA3, and to any subsequent releases until otherwise indicated in new editions or technical newsletters.

The changes for this edition are summarized under "Summary of Changes" following the table of contents. Specific changes are indicated by a vertical bar to the left of the change. A vertical bar to the left of a figure caption indicates that the figure has changed. Editorial changes that have no technical significance are not noted.

Changes are made periodically to this publication; before using this publication in connection with the operation of IBM systems, consult the latest *IBM System/370, 30xx, 4300, and 9370 Processors Bibliography*, GC20-0001, for the editions that are applicable and current.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM licensed program in this publication is not intended to state or imply that only IBM's program may be used. Any functionally equivalent program may be used instead.

Requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality. If you request publications from the address given below, your order will be delayed because publications are not stocked there.

A Reader's Comment Form is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Department J57, P.O.Box 49023, San Jose, California, U.S.A. 95161-9023. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Trademarks

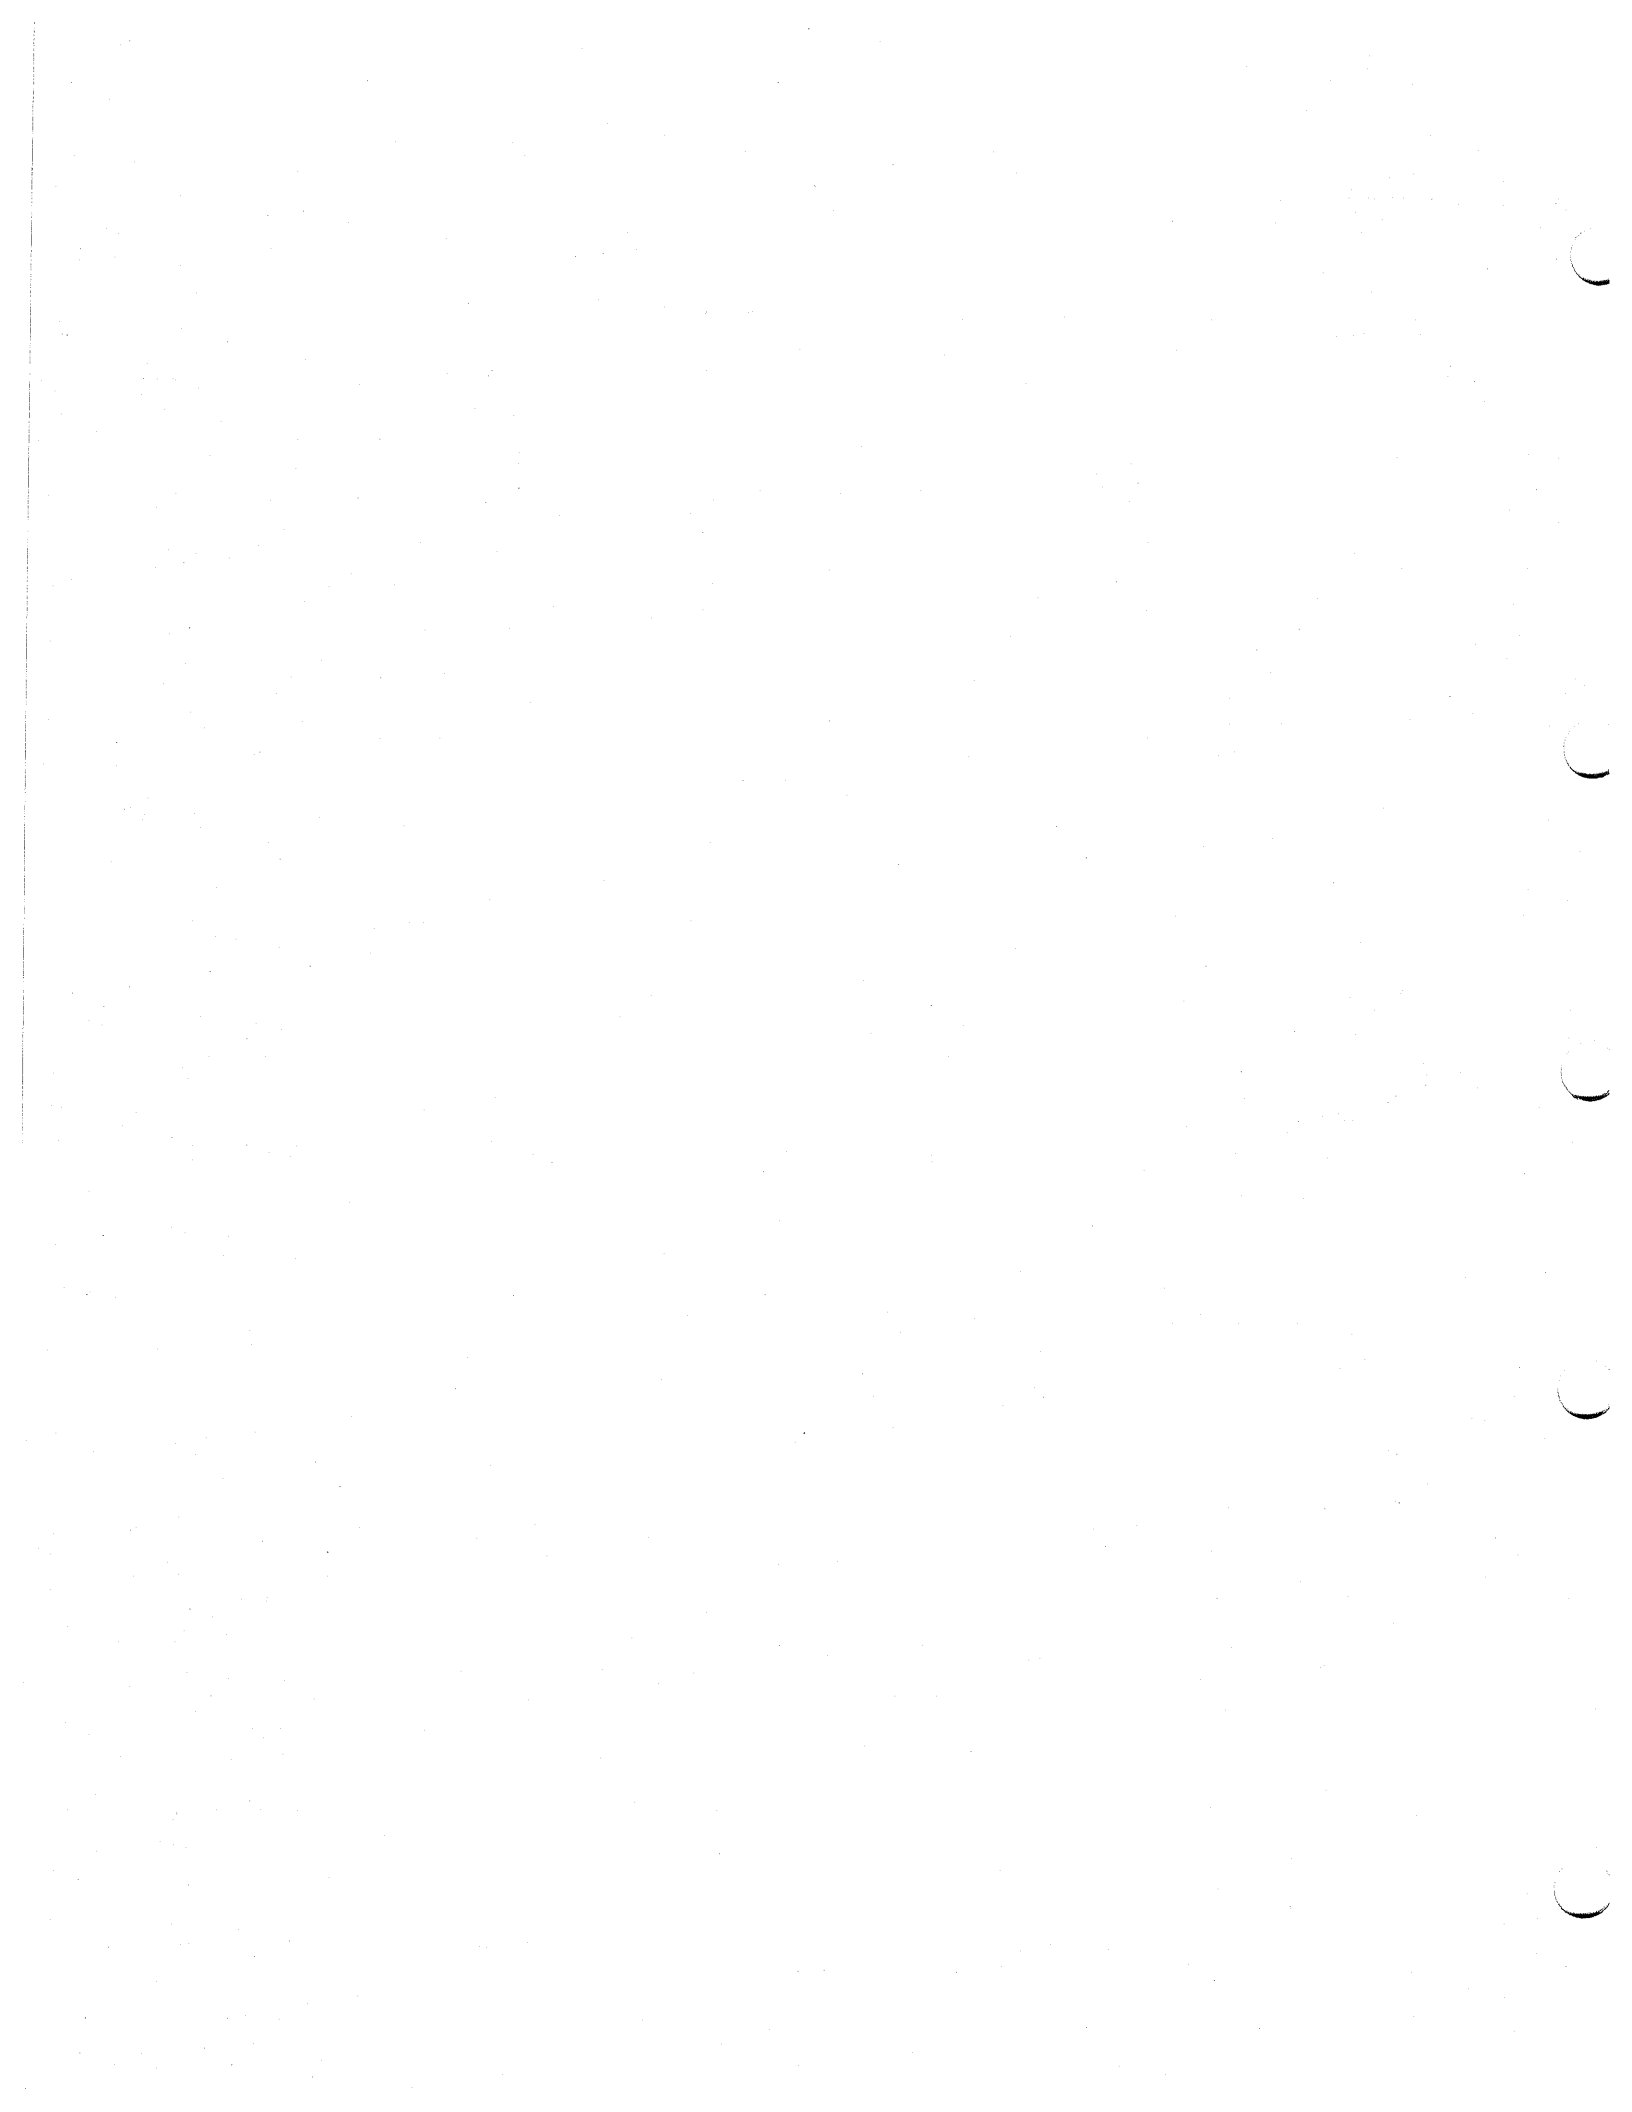The following names have been adopted by IBM for trademark use and are used throughout this publication:

MVS/DFP™

MVS/ESA™

MVS/SP™

MVS/XA™

Contents

Summary of Changes

Preface

About This Book

This book is intended to help you manage storage under MVS/DFP. It contains explanations of how to initialize the MVS/DFP **Storage Management Subsystem** (SMS), maintain it, and perform storage management tasks with the help of **Interactive Storage Management Facility** (ISMF) panels. The information in this book must not be used for programming purposes.

The **Storage Management Library** (SML) provides additional information about SMS, including suggestions and recommendations for automating your storage management.

To get the most out of this manual, you should have access to the applications and functions available through the ISMF Primary Option Menu for Storage Administrators. If you want to use the applications and functions available through the ISMF Primary Option Menu for End Users, see the *MVS/ESA Interactive Storage Management Facility User's Guide*.

Required Product Knowledge

To use this book effectively, you should be familiar with the concepts presented in the SML. The SML describes storage management concepts and provides recommendations for the transition toward system-managed storage. This manual assumes you have begun preparing for system-managed storage by implementing some of the recommendations in the SML.

You should also be familiar with the storage management functions provided by Data Facility Data Set Services (DFDSS) and Data Facility Hierarchical Storage Manager (DFHSM). DFDSS moves data from one device to another, backs up and recovers data sets, and reduces free-space fragmentation on DASD volumes. DFHSM provides automatic space management and availability functions.through a hierarchy of storage devices.

You should also be familiar with the Device Support Facility (DSF).

Required Publications

You should be familiar with the information presented in the following publications:

Publication Title	Order Number
MVS Storage Management Library: Storage Management Reader's Guide	GC26-4403
MVS Storage Management Library: Focus on Storage Management	GC26-4404

Publication Title	Order Number
MVS Storage Management Library: Leading an Effective Storage Administration Group	SC26-4405
MVS Storage Management Library: Storage Management Subsystem Migration Planning Guide	SC26-4406
MVS Storage Management Library: Managing Storage Pools	SC26-4407
MVS Storage Management Library: Managing Data Sets	SC26-4408
MVS Storage Management Library: Configuring Storage Subsystems	SC26-4409

You can order the entire Storage Management Library in a three-ring binder using order number SBOF-1241.

Related Publications

Some publications from the MVS/SP Version 3 library are referenced in this book. The *MVS/ESA Library Guide for System Product Version 3*, GC28-1563, contains a complete listing of the MVS/SP Version 3 publications and their counterparts for the prior version.

The *MVS/ESA Data Facility Product Version 3: Master Index*, GC26-4512, contains both an index to the MVS/DFP library and a summary of the changes made to the library. You can use it to:

* Find information in other MVS/DFP Publications
* Determine how new programming support changes information in the MVS/DFP library
* Determine which MVS/DFP publications have been changed.

Referenced Publications

Within the text, references are made to the publications listed below:

Short Title	Publication Title	Order Number
DFDSS: User's Guide	*Data Facility Data Set Services: User's Guide*	SC26-4388
ICKDSF User's Guide and Reference	*Device Support Facilities User's Guide and Reference*	GC35-0033
DFHSM: General Information	*Data Facility Hierarchical Storage Manager: Version 2 Release 4 General Information*	GC35-0092
Cache Device Administration	*Cache Device Administration*	GC35-0101
DFP: Customization	*MVS/ESA Data Facility Product Version 3: Customization*	SC26-4504

Short Title	Publication Title	Order Number
DFP: General Information	*MVS/ESA Data Facility Product Version 3: General Information*	GC26-4507
Access Method Services Reference	*MVS/ESA Integrated Catalog Adminstration: Access Method Services Reference*	SC26-4500
Access Method Services Reference	*MVS/ESA Integrated Catalog Adminstration: Access Method Services Reference*	SC26-4500
ISMF User's Guide	*MVS/ESA Interactive Storage Management Facility User's Guide*	SC26-4508
JCL Reference	*MVS/ESA JCL Reference*	GC28-1829
System Messages Volume 1	*MVS/ESA Message Library: System Messages Volume 1*	GC28-1812
System Messages Volume 2	*MVS/ESA Message Library: System Messages Volume 2*	GC28-1813
Operations: System Commands	*MVS/ESA Operations: System Commands*	GC28-1826
RACF Security Administrator's Guide	*Resource Access Control Facility (RACF) Security Administrator's Guide*	
System Generation	*MVS/ESA System Generation*	GC28-1825
Initialization and Tuning	*MVS/ESA System Programming Library: Initialization and Tuning*	GC28-1828
VSAM Administration Guide	*MVS/ESA VSAM Administration Guide*	SC26-4518

Note on Examples

For purposes of illustration, the examples in this manual use the base configuration, SMS classes and storage groups, and ACS routines that are part of the Starter Set, which is described in Appendix B, "Starter Set" on page 143.
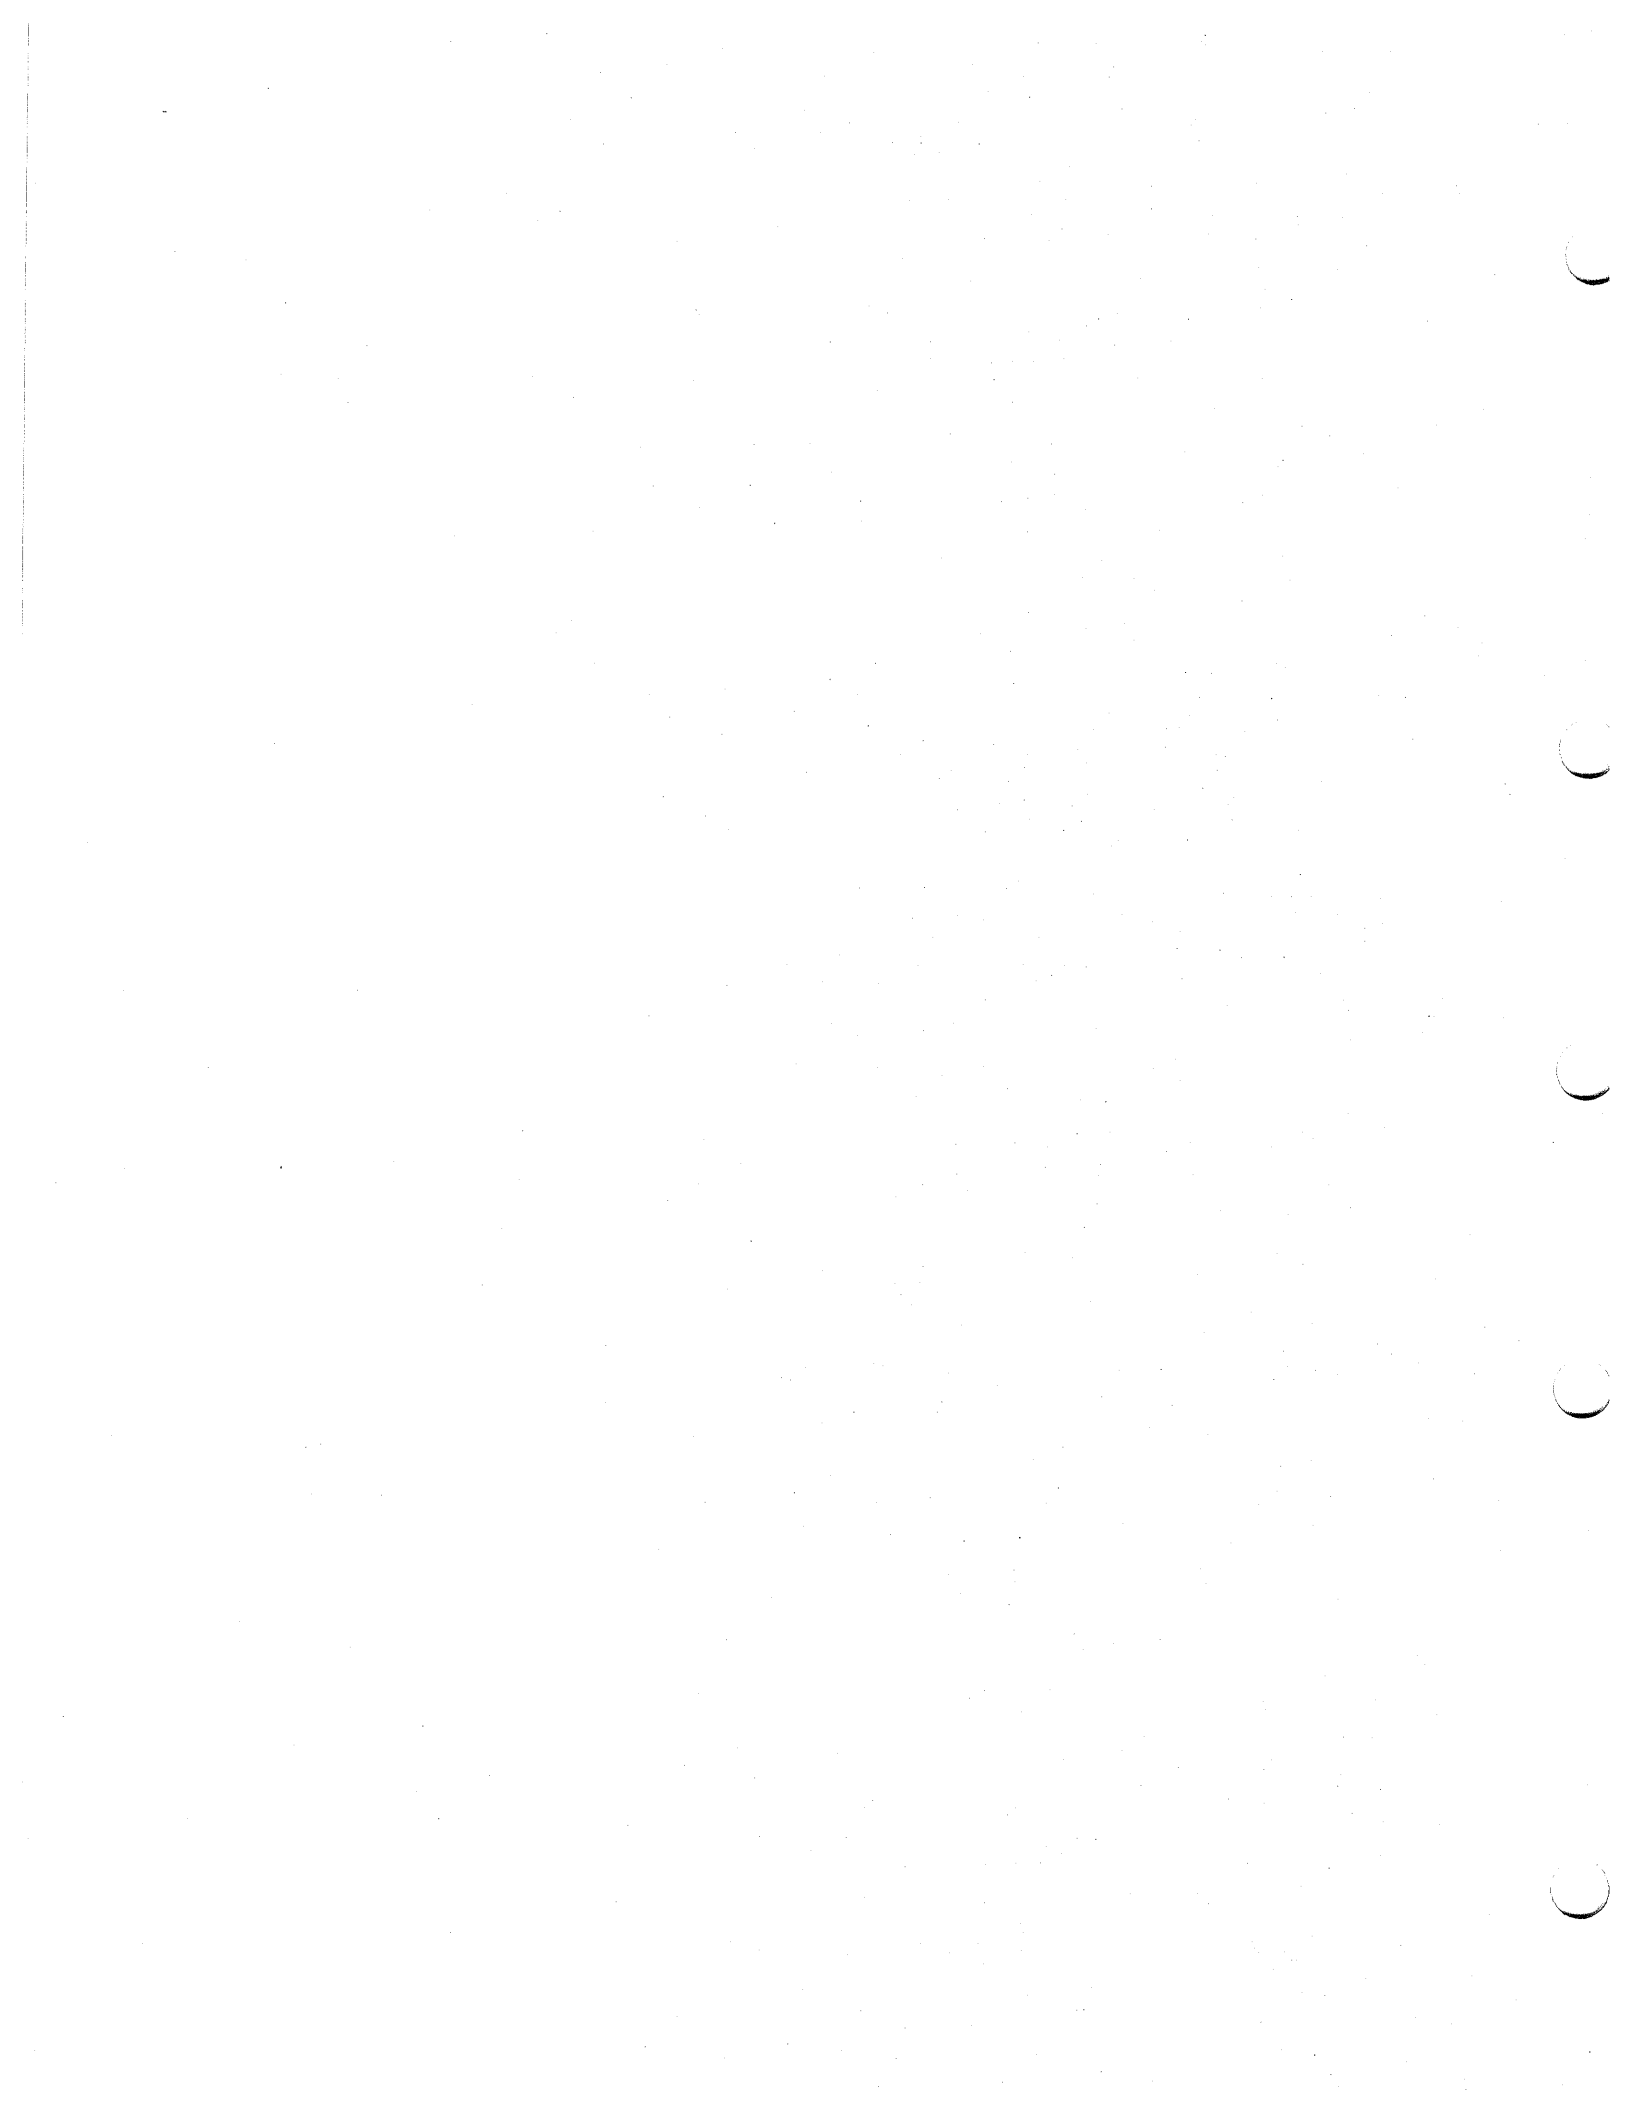
Chapter 1. Introducing the Storage Management Subsystem

MVS/ESA Data Facility Product Version 3 Release 1 introduces the **Storage Management Subsystem** (SMS). SMS provides new functions for data and space management. An important step toward system-managed storage, SMS:

- Facilitates the management of storage growth
- Improves the use of storage space
- Reduces the effort of device conversion and coexistence
- Provides centralized control of external storage
- Exploits the capabilities of available hardware.

SMS manages an installation's storage according to the currently active storage management policy. Through the Interactive Storage Management Facility (ISMF), you define an installation storage management policy in an **SMS configuration**. An SMS configuration contains the following:

- CDS base configuration
- Constructs
- Automatic class selection (ACS) routines.

The **Control Data Set (CDS) base configuration** identifies the systems that the SMS configuration is to manage. These systems constitute an **SMS complex**. The CDS base configuration also contains installation defaults.

Constructs are lists of traits and characteristics that are assigned to data sets and volumes. An SMS configuration can contain the following four types of constructs:

Storage group allows you to define a list of volumes and manage them as if they were one large, single volume. SMS applies the properties you assign to a storage group to all the volumes within the storage group.

Management class allows you to define different levels of migration, backup, and retention services. Through management class, you can associate a level of service with a data set that is independent of the volume containing the data set.

Storage class allows you to define different levels of performance and availability services. Through storage class, you can separate the level of service for a data set from physical device characteristics.

Data class allows you to define allocation defaults. Through data class, you can simplify and standardize the allocation of new data sets.

An SMS configuration can contain multiple constructs of each type. Figure 1 on page 2 shows how the SMS classes and storage groups are related to a data set.
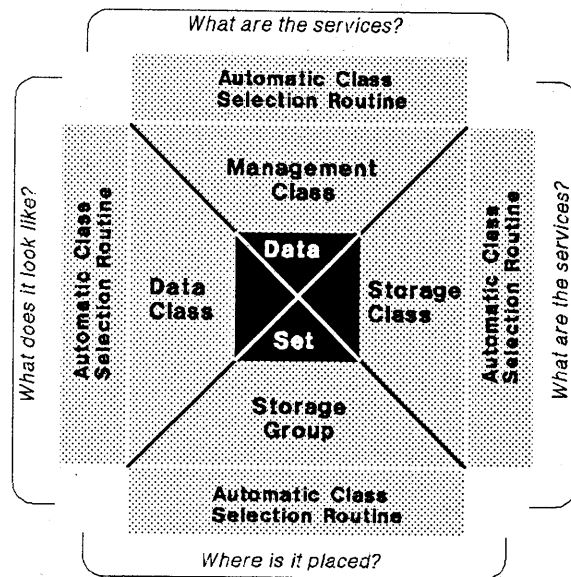
Figure 1. SMS and its relationship to a data set

ACS routines determine the SMS classes and storage groups for data sets. You can also use ACS routines to control the transition of data sets to and from SMS management.

The chapters in this manual provide greater detail on these topics. See *MVS/ESA Data Facility Product Version 3: General Information* for a more detailed introduction to the Storage Management Subsystem.

Using the Interactive Storage Management Facility

ISMF provides a series of applications for storage administrators to define and manage SMS configurations. You can use these applications to:

- Define SMS CDS base configuration information

- Define, alter, delete, or copy individual SMS classes and storage groups

- Display parameters and values of individual SMS classes and storage groups

- Generate lists of SMS classes and storage groups

- Edit ACS routines

- Define, alter, and execute ACS test cases

- Validate the correctness and completeness of an SMS configuration

- Activate an SMS configuration.

- Display, define, alter, or delete storage group information pertaining to specific volumes.

You can select ISMF from the Interactive Systems Productivity Facility (ISPF) primary option menu. Figure 2 on page 3 shows the ISMF Primary Option Menu for Storage Administrators.

```
                    ISMF PRIMARY OPTION MENU
ENTER SELECTION OR COMMAND ===>

SELECT ONE OF THE FOLLOWING:

0   ISMF PROFILE            - Change ISMF user profile
1   DATA SET                - Perform Functions Against Data Sets
2   VOLUME                  - Perform Functions Against Volumes
3   MANAGEMENT CLASS        - Specify Data Set Backup and Migration Criteria
4   DATA CLASS              - Specify Data Set Allocation Parameters
5   STORAGE CLASS           - Specify Data Set Performance and Availability
6   STORAGE GROUP           - Specify Volume Names and Free Space Thresholds
7   AUTOMATIC CLASS SELECTION - Specify ACS Routines and Test Criteria
8   CONTROL DATA SET        - Specify System Names and Default Criteria
X   EXIT                    - Terminate ISMF




USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 2. ISMF Primary Option Menu for Storage Administrators

This primary option menu differs from the one that your end users see. The
ISMF Primary Option Menu for End Users does not have the **Storage Group**,
Automatic Class Selection, and **Control Data Set** options that are available to
you. Additionally, the management class, data class, and storage class appli-
cations available through the ISMF Primary Option Menu for Storage Adminis-
trators allow you to define, alter, copy, and delete SMS classes. End users can
only list the available SMS classes and display the attributes of individual SMS
classes.

Installing Required and Corequisite Programs and Products

To use MVS/DFP Version 3, you must install the supporting levels of ISPF and
MVS/SP JES2 on your system. To make full use of SMS functions, ISMF com-
mands, and ISPF edit and browse functions, you need to install the supporting
levels of the following corequisite products or their functional equivalents:

- TSO/Extensions (TSO/E)
- Resource Measurement Facility (RMF)
- Data Facility Hierarchical Storage Manager (DFHSM)
- Data Facility Data Set Services (DFDSS)
- Resource Access Control Facility (RACF)
- Device Support Facility (ICKDSF)
- Environmental Recording, Editing, Printing program (EREP)
- Interactive System Productivity Facility/Program Development Facility
  (ISPF/PDF).
- IBM 3990 Storage Control Model 3

This manual assumes that you have installed all of these products or their func-
tional equivalents. For purposes of consistency and illustration, the text and
examples throughout this manual make reference only to the functions and
capabilities of the products listed above. For information on the required level

of these products, see *MVS/ESA Data Facility Product Version 3: General Information.*

Implementing SMS

Before you implement a storage management policy by activating an SMS configuration, you need to follow the steps outlined in Figure 3 on page 5. The chapters in this manual that describe these steps appear to the right of the respective boxes.

When preparing for SMS, you need to allocate **control data sets**, which will contain information used by SMS. Next, you need to modify and create SYS1.PARMLIB members to identify SMS to all the systems in the SMS complex. As a final preparation step, you need to establish access to the ISMF Primary Option Menu for Storage Administrators, which is shown in Figure 2 on page 3.

Then you define the base configuration, which identifies the systems within the SMS complex. After defining a base configuration, you can define the SMS classes and storage groups that you want SMS to assign to your data sets. Then you can define ACS routines to determine the SMS classes and storage groups.

After testing the ACS routines, you can validate them individually to check for errors. You should then validate the entire SMS configuration to check for errors that exist among its related parts. Finally, you can activate the valid SMS configuration.

```
┌─────────────────────────────────────────┐
│  1. Allocate control data sets          │
│                                         │
│  2. Modify and create SYS1.PARMLIB members │   Chapter 2
│                                         │
│  3. Access storage administrator ISMF options │
└─────────────────────────────────────────┘
                     │
                     ▼
┌─────────────────────────────────────────┐
│  4. Define base configuration           │   Chapter 3
│                                         │
│  5. Define SMS classes and storage groups │   Chapters 4 - 7
│                                         │
│  6. Define and test ACS routines,       │   Chapter 8
│                                         │
│  7. Validate ACS routines and           │
│        the SMS configuration            │
└─────────────────────────────────────────┘
                     │
                     ▼
┌─────────────────────────────────────────┐
│  7. Activate SMS configuration          │   Chapter 9
└─────────────────────────────────────────┘
```

Figure 3. Activating an SMS Configuration
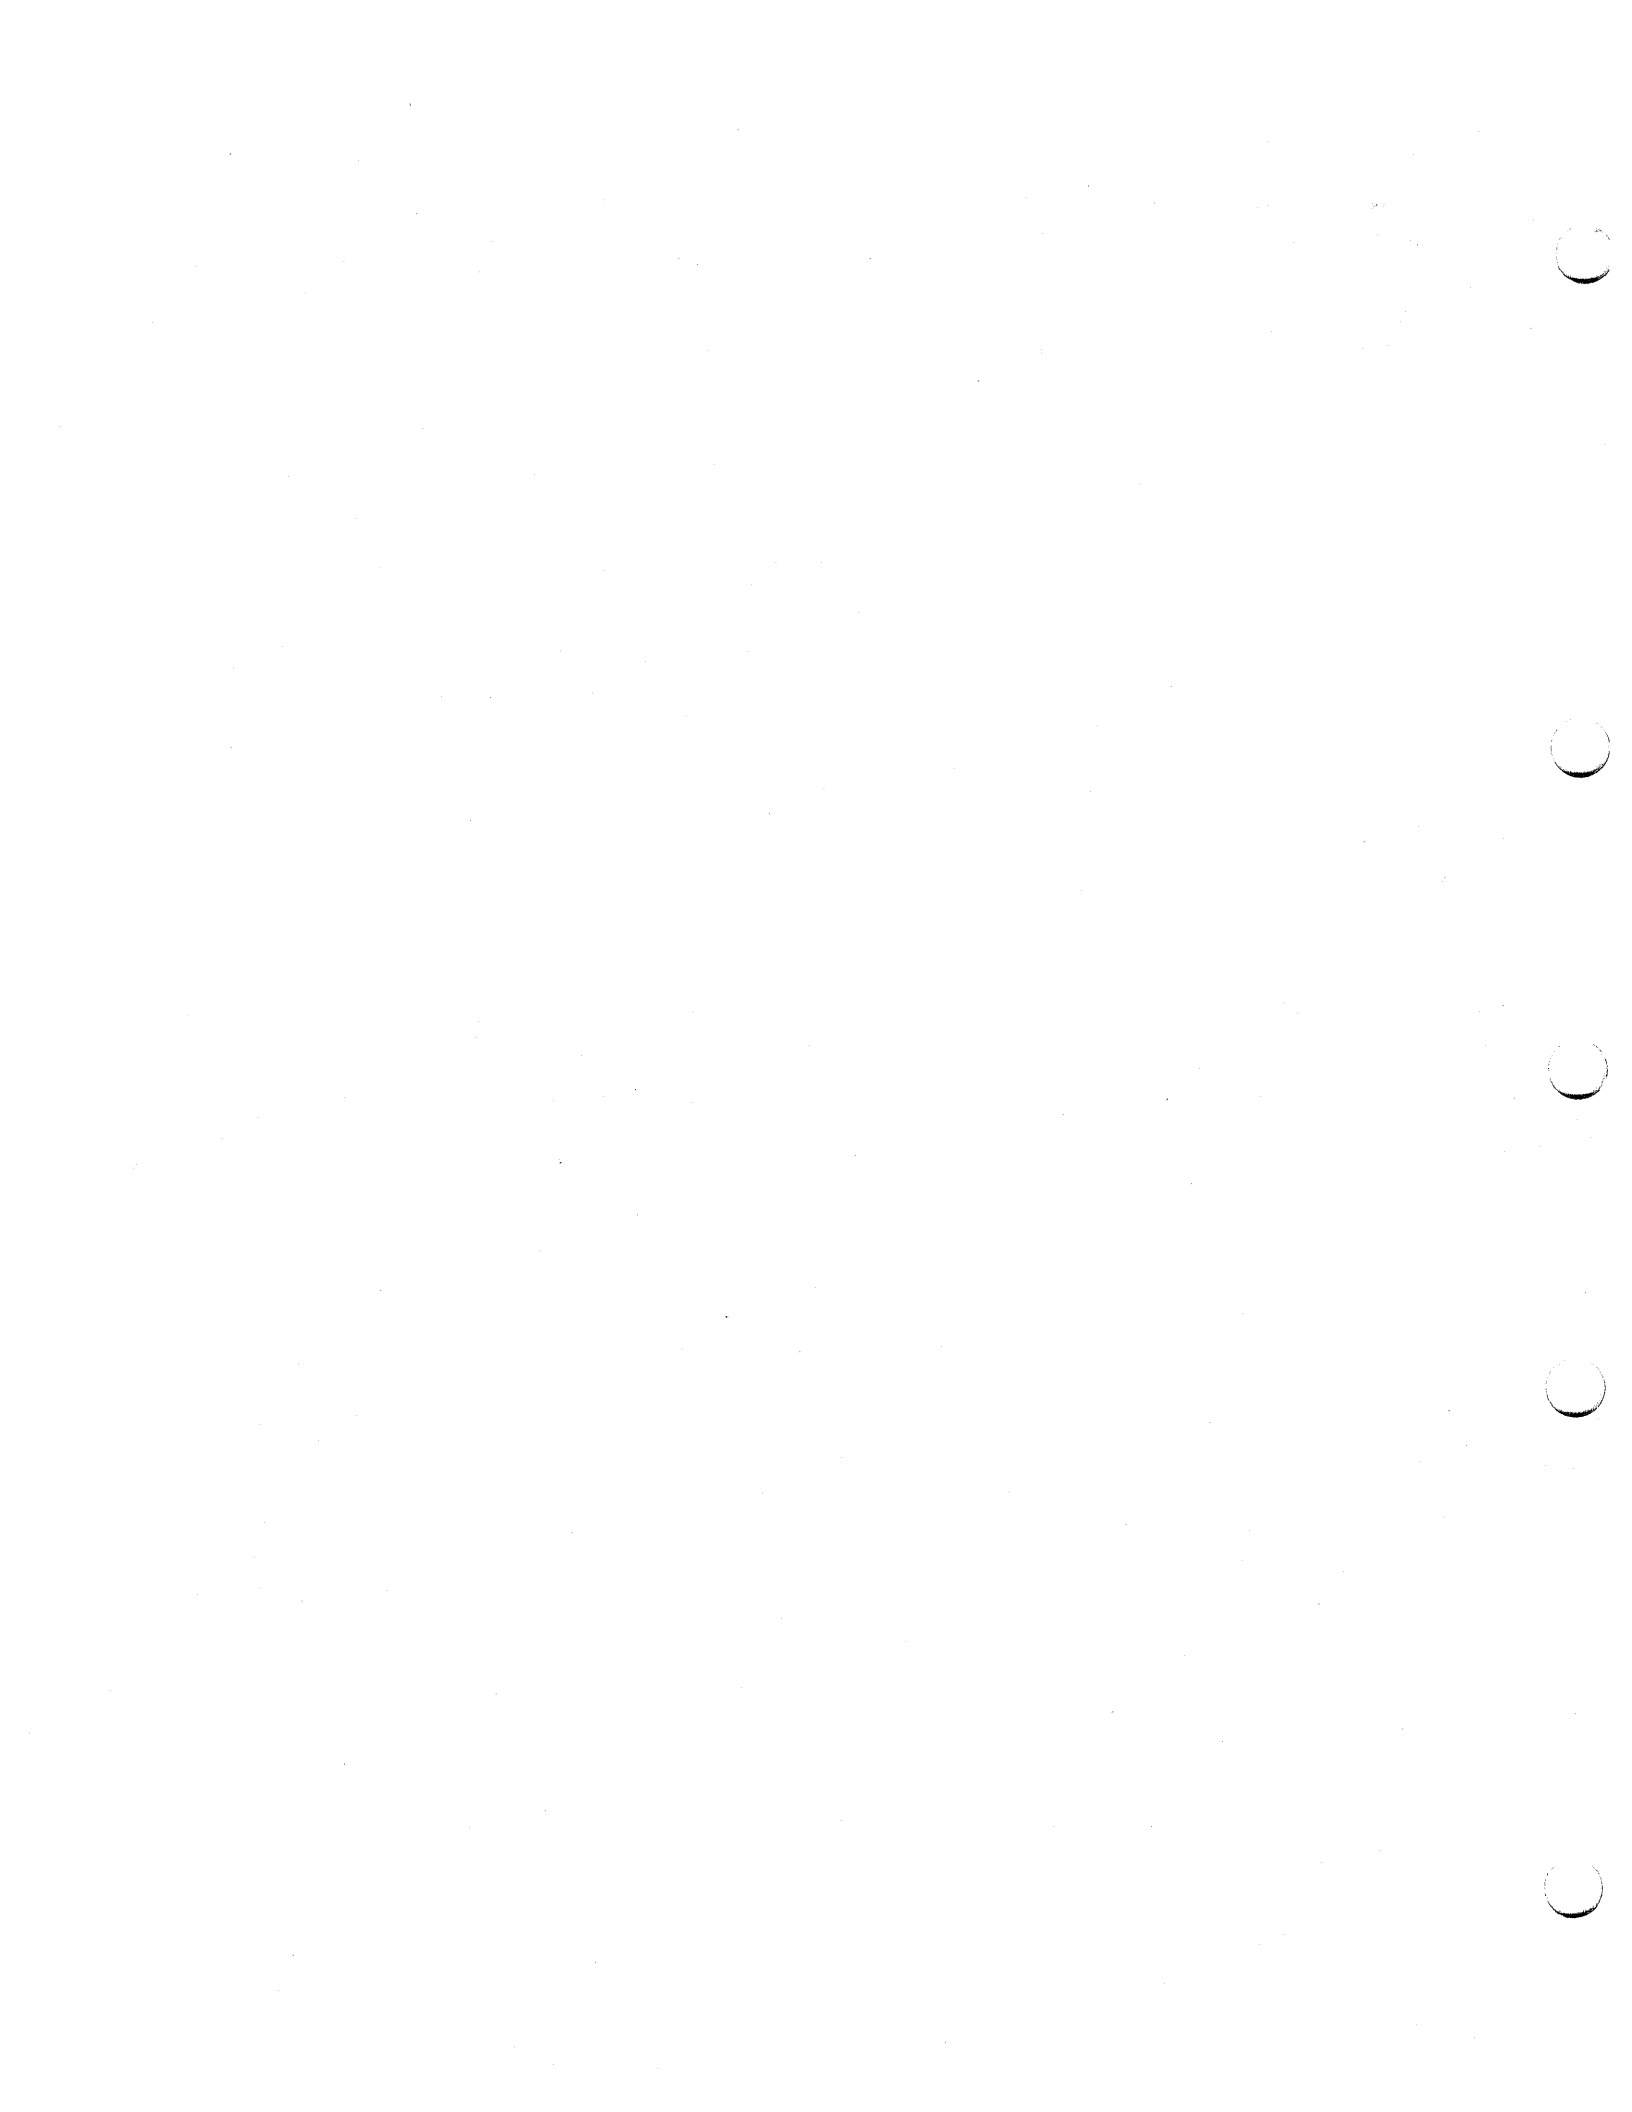
Chapter 2. Preparing for SMS

Before defining and activating an SMS configuration, you need to perform the following preparatory steps:

- Allocate control data sets to contain your SMS configuration and to permit the systems in your complex to communicate with each other.

- Modify SYS1.PARMLIB, which contains two members that direct the initialization and activation of SMS.

- Access the ISMF Primary Option Menu for Storage Administrators (shown in Figure 2 on page 3).

This chapter describes how to perform these preliminary steps so that you can begin defining a base configuration for an SMS configuration.

For additional information regarding the planning and preparation for SMS, see *MVS SML: Storage Management Subsystem Migration Planning Guide.*

Allocating Control Data Sets

Before you can activate an SMS configuration, you need to allocate control data sets and define their contents. Control data sets are virtual storage access method (VSAM) linear data sets (LDS) that contain base configuration information, SMS class and storage group definitions, and ACS routines.

Before SMS is active, you can allocate control data sets by executing access method services in batch mode or issuing TSO/E commands (explained in this section). You define and alter the contents of control data sets using ISMF. SMS uses three types of control data sets: the **source control data set** (SCDS), the **active control data set** (ACDS), and the **communications data set** (COMMDS).

Source Control Data Set

An SCDS contains an SMS configuration, which defines a storage management policy. You can define any number of SMS configurations each of which is its own SCDS. Then, you select one SMS configuration to be the installation storage management policy and make an active working copy of it in an ACDS.

Active Control Data Set

When you activate an SCDS, its contents are copied to an ACDS. An ACDS contains a copy of the most recently activated configuration. All systems in the SMS complex use this configuration to manage storage. You can define any number of ACDS, but only one contains the currently active storage management policy. "Initializing SMS through IGDSMSmm" on page 11 explains how to specify which ACDS contains the current storage management policy.

You can modify the SCDS from which your current storage management policy was activated without disrupting operations, because SMS manages storage with a copy of the SMS configuration (an ACDS) rather than with the original (an SCDS). While SMS manages storage using an ACDS, you can:

- Create a backup copy of the SCDS
- Build a new SCDS

- Update the SCDS from which the ACDS was activated
- Modify any SCDS.

Figure 4 shows the relationship among the SCDS and ACDS in an installation.
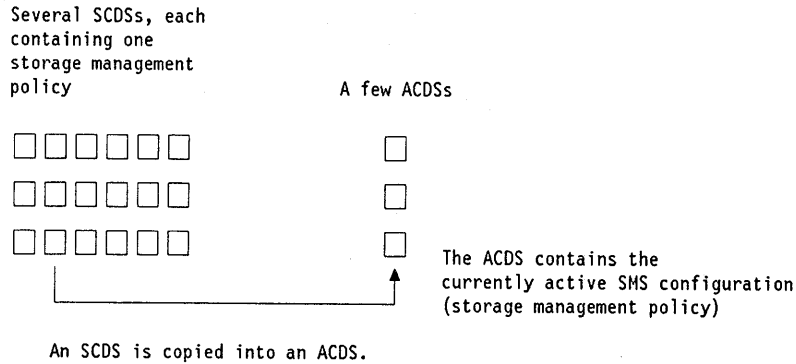
Figure 4. Relationship among SCDSs and ACDSs in an Installation

Communications Data Set

The COMMDS serves as the primary means of communication among systems in the SMS complex. The active systems in an SMS complex access the COMMDS for current SMS complex information. An SMS complex can contain up to eight systems.

The COMMDS contains the name of the ACDS containing the currently active storage management policy, the current utilization statistics for each SMS-managed volume, and other system information. You can define any number of COMMDS, but only one can be active in an SMS complex.

Calculating the Size of Control Data Sets

Before you allocate control data sets, you need to estimate their size. When calculating the size of either an ACDS or an SCDS, you have to account for system and base configuration information, SMS class and storage group definitions, and ACS routines.

The following formula summarizes these three pieces of information to help you determine the size of a control data set.

DC = Estimated number of data classes
MC = Estimated number of management classes
SC = Estimated number of storage classes
SG = Estimated number of storage groups
VOL = Estimated number of DASD volumes in the SMS complex to be SMS managed

SCDS/ACDS
size in bytes = 123000 + (1600 * (SG + MC + SC + DC)) + (800 * VOL)

The 123000 bytes represent fixed data fields. As an example, if you have:

Five storage groups
Four management classes
Two storage classes
Five data classes

Four ACS routines
40 DASD volumes

the equation yields a value of **180,600 bytes** for control data set allocation, which you can round to 200K (where 1K = 1024 bytes).

When calculating the size of a COMMDS, you have to account for both system and volume information. The following formula summarizes the two pieces of information to help you determine the size of a COMMDS:

COMMDS size (bytes) = 8192 + (28 * VOL)

As an example, if you have 40 DASD volumes in the SMS complex, you need to allocate 9,312 bytes for the COMMDS, which you can round to 10K.

Allocating an SCDS

To allocate the 200K SCDS described in the previous section, you can execute access method services using the following JCL stream:

```
//ALLOC  JOB  MSGCLASS=A,USER=STGADM1,NOTIFY=STGADM1,
//            MSGLEVEL=(1,1)
//STEP  EXEC PGM=IDCAMS
//SYSUDUMP DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
  DEFINE CLUSTER(NAME(SMS.SCDS1.SCDS) LINEAR VOL(SMSV01) -
  KILOBYTES(200) SHAREOPTIONS(2,3)) -
  DATA(NAME(SMS.SCDS1.SCDS.DATA))

/*
```

This JCL generates a VSAM LDS named SMS.SCDS1.SCDS. You can combine the DEFINE commands for all your allocations into the same JCL job stream, but this example shows only one for purposes of illustration. After allocating an SCDS, you define its contents through ISMF dialogs.

You should allocate an SCDS on a device shared by all systems in the SMS complex. If you allocate an SCDS on a device that is not shared by all the systems, then you can only activate the SCDS from systems that have access to it.

Allocating an ACDS

To allocate a 200K ACDS for the SCDS that you created in the previous allocation, execute access method services using the following JCL stream:

Note: When allocating an ACDS, the share options must be different from the share options you specified in the SCDS.

```
//ALLOC  JOB  MSGCLASS=A,USER=STGADM1,NOTIFY=STGADM1,
//            MSGLEVEL=(1,1)
//STEP  EXEC PGM=IDCAMS
//SYSUDUMP DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
  DEFINE CLUSTER(NAME(SMS.ACDS1.ACDS) LINEAR VOL(SMSV02) -
  KILOBYTES(200) SHAREOPTIONS(3,3)) -
  DATA(NAME(SMS.ACDS1.ACDS.DATA))

/*
```

This JCL generates a VSAM LDS named SMS.ACDS1.ACDS. You can combine
the DEFINE commands for all your allocations into the same JCL job stream,
but this example shows only one for purposes of illustration.

An ACDS must reside on a shared volume, accessible from all systems in the
SMS complex. To ease recovery in case of failure, the ACDS should reside on
a different volume than the COMMDS. Also, you should allocate a spare ACDS
on a different shared volume. Chapter 11, "Recovering SMS Information" on
page 105 provides additional information on the backup and recovery of control
data sets.

You create the contents of an ACDS by activating a valid SCDS. The distinction
between a valid and an invalid SCDS is described in "Defining SCDS Base Con-
figuration Information" on page 18.

Allocating a COMMDS

To allocate the 10K COMMDS described earlier, execute access method ser-
vices using the following JCL stream:

```
//ALLOC JOB  MSGCLASS=A,USER=STGADM1,NOTIFY=STGADM1,
//           MSGLEVEL=(1,1)
//STEP  EXEC PGM=IDCAMS
//SYSUDUMP DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *

  DEFINE CLUSTER(NAME(SMS.COMMDS1.COMMDS) LINEAR VOL(SMSVOL) -
  KILOBYTES(10) SHAREOPTIONS(3,3)) -
  DATA(NAME(SMS.COMMDS1.COMMDS.DATA))

/*
```

This JCL generates a VSAM LDS SMS.COMMDS1.COMMDS.

The COMMDS must reside on a shared volume, accessible from all systems in
the SMS complex. To ease recovery in case of failure, the COMMDS should
reside on a different volume than the ACDS. Also, you should allocate a spare
COMMDS on a different shared volume. Chapter 11, "Recovering SMS
Information" on page 105 provides additional information on the backup and
recovery of control data sets.

Modifying SYS1.PARMLIB

The IGDSMSmm and IEFSSNnn members of SYS1.PARMLIB direct the initializa-
tion and activation of SMS. IGDSMSmm is a new SYS1.PARMLIB member that
provides initialization parameters to SMS. IEFSSNnn identifies SMS to MVS.
This section explains how to create an IGDSMSmm member in SYS1.PARMLIB
and provides two methods to define SMS to MVS through IEFSSNnn.

Initializing SMS through IGDSMSmm

For each system in the SMS complex, you must build an IGDSMSmm member in SYS1.PARMLIB. IGDSMSmm is a new PARMLIB member that contains SMS initialization control information. IGD is the SMS component identifier, and SMS identifies the member as a repository of SMS initialization control information. 'mm' is any value allowed by the naming conventions for SYS1.PARMLIB members. It has a default value of 00.

Every system must have IGDSMSmm member in SYS1.PARMLIB that specifies a required ACDS/COMMDS pair. The data sets that you specify for this ACDS/COMMDS pair must be the same for every system in an SMS complex. Whenever you change the ACDS or COMMDS, you must update the IGDSMSmm member of SYS1.PARMLIB for every system in the SMS complex so that it specifies the same data sets.

The IGDSMSmm member contains the following parameters, which direct the SMS initialization process and specify the name of the data set containing SMS class and storage group definitions:

```
SMS ACDS(dsname) COMMDS(dsname)
    [INTERVAL(nnn)]
    [DINTERVAL(nnn)]
    [REVERIFY(YES|NO)]
    [ACSDEFAULTS(YES|NO)]
    [TRACE(OFF|ON)]
    [SIZE(nnn|nnnK|nnnM)]
    [TYPE(ALL|ERROR)]
    [JOBNAME(jobname|*)]
    [ASID(asid|*)]
    [SELECT(event, event, ...|ALL)]
    [DESELECT(event, event, ...)]
```

Note: SMS, ACDS, COMMDS, REVERIFY, and ACSDEFAULTS need to be the same for all your systems. The rest of the parameters can be different.

SMS

> Identifies the member as a repository of SMS initialization control information. **SMS** is a required keyword.

ACDS(dsname)

> dsname identifies the name of the data set containing the configuration. **ACDS** is a required keyword. If you omit dsname, the operator will be prompted for a value.

COMMDS(dsname)

> dsname represents the name of the COMMDS. **COMMDS** is a required keyword. If you omit dsname, the operator will be prompted for a value.

INTERVAL(nnn)

> nnn represents the **synchronization time interval** of the system, which is the number of seconds that SMS allows before this system checks the COMMDS for news from other systems in the SMS

complex. The **INTERVAL** keyword is optional. You can specify a value from 1 to 999, and the default is 15.

DINTERVAL(*nnn***)**

> *nnn* represents the number of seconds SMS allows to elapse before it reads device statistics. The **DINTERVAL** parameter is optional and applies only to the volumes behind an IBM 3990 Storage Control Model 3 or later model. The default is 150.

REVERIFY(YES|NO)

> A value of **NO** instructs SMS to verify a user's authority to allocate a new data set, to use a storage class, and to use a management class at job interpretation time only. A value of **YES** instructs SMS to perform the verification at both job interpretation and job execution time. JES2 environments ignore the **REVERIFY** keyword. The **REVERIFY** keyword is optional. It has a default value of **NO**.

ACSDEFAULTS(YES|NO)

> The **ACSDEFAULTS** keyword indicates whether SMS is to initialize the following ACS routine variables from an additional call to RACF:

> - &APPLIC
> - &DEF_DATACLAS
> - &DEF_MGMTCLAS
> - &DEF_STORCLAS

> If you specify **NO**, these variables will have no values associated with them. The **ACSDEFAULTS** keyword is optional. It has a default value of **NO**.

TRACE(OFF|ON)

> The **TRACE** keyword indicates whether you want to activate the SMS trace facility. The default is **ON**. The SMS trace facility records trace records in the SMS address space. Refer to the *DFP: Diagnosis Reference* for additional information.

SIZE(*nnn***|***nnn***K|***nnn***M)**

> The **SIZE** keyword specifies the size of the SMS trace table in bytes. When the unit is **K**, value can range from 0K to 255,000K, and it is rounded up to the nearest 4K unit. When the unit is **M**, value can range from 0M to 255M. If you specify a value of 0, no tracing is performed. The default is 128K.

TYPE(ALL|ERROR)

> The **TYPE** keyword specifies whether you want to trace all events or only errors. The default is **ERROR**.

JOBNAME(*jobname***|*)**

> The **JOBNAME** keyword limits tracing to a certain job or permits tracing on all jobs. The default is to trace all jobs, *.

ASID(*asid***|*)**

> The **ASID** keyword limits tracing to a certain address space or permits it for all address spaces. The default is to trace all address spaces, *.

SELECT(*event, event, ...***|ALL)**

> The **SELECT** keyword adds items to the list of events to be traced. The default is **ALL**.

DESELECT(*event, event, ...*)

The **DESELECT** keyword deletes items from the list of events to be traced. It has no defaults. If you specify SELECT and DESELECT keyword events that conflict, the keyword that appears last has final authority.

The events that you can specify on SELECT and DESELECT are:

MODULE	SMS module entry/exit
SMSSJF	SMS/SJF interfaces
SMSSSI	SMS/SSI interfaces
ACSINT	ACS services interfaces
OPCMD	Operator commands
CONFC	Configuration change
CDSC	Control data set changes
CONFS	SMS Configuration services
MSG	SMS Message services
ERR	SMS Error recovery and recording services
CONFR	Return data from an active configuration
CONFA	Activate a new configuration
ACSPRO	Perform ACS processing
IDAX	SMS interpreter/dynamic allocation
DISP	SMS disposition processing exit
CATG	SMS catalog services
VOLREF	SMS VOLREF services
SCHEDP	SMS scheduling services (prelocate catalog orientation)
SCHEDS	SMS scheduling services (system select)
VTOCL	SMS VTOC/data set services (allocate existing data set)
VTOCD	SMS VTOC/data set services (delete existing data set)
VTOCR	SMS VTOC/data set services (rename existing data set)
VTOCC	SMS VTOC/data set services (create new data set)
VTOCA	SMS VTOC/data set services (add a volume to a data set)
RCD	SMS Recording services
ALL	All of the above options

For more information on the parameters of the IGDSMSmm member, see *MVS/ESA System Programming Library: Initialization and Tuning.*

Defining SMS to MVS through IEFSSNnn

Before you can activate SMS, MVS must recognize SMS as a valid subsystem. You can define SMS to MVS in one of two ways. The first method involves placing a record for SMS in the existing IEFSSNnn member that is used for IPL. The record for SMS should appear before any other records. The format for the SMS record is shown in "Writing the SMS Record" on page 14.

The second method involves creating a new IEFSSNnn member that only contains the record for SMS. To identify the IEFSSNnn member containing the SMS record, you must update the IEASYSyy member that you use to IPL. In the SSN parameter of IEASYSyy, the suffix identifying the IEFSSNnn member that contains the SMS record should appear first. For example, in:

SSN=(nn,01,...)

'nn' identifies IEFSSNnn as the member containing the SMS definition (the SMS record). '01' identifies IEFSSN01 as another member defining any other subsystem. You can specify additional members to define more subsystems such as:

SSN=(nn,01,...)

at IPL time.

Writing the SMS Record

To define SMS to MVS, you must place a record for SMS in an IEFSSNnn member. The format of the SMS record is shown in Figure 5.

Figure 5. Format of the SMS Record in the IEFSSNnn member of SYS1.PARMLIB

SMS   Within the record, SMS identifies and defines the Storage Management Subsystem to MVS.

IGDSSIIN   IGDSSIIN is the subsystem initialization routine for SMS. Including it in the SMS record automatically starts SMS during every future IPL. Omitting it from the SMS record defines SMS as a valid subsystem to MVS, but SMS does not automatically start during future IPLs.

ID   The ID field identifies a repository for SMS initialization information. The two characters (mm) represent the suffix for the SMS initialization member, IGDSMSmm. If you omit the ID field, the system uses the IGDSMS00 member as a default.

PROMPT   The PROMPT parameter indicates how much control you want the operator to have during the rest of SMS initialization. DISPLAY requests that the contents of the IGDSMSmm member be displayed, but you cannot modify them. YES specifies that a Write to Operator with Reply (WTOR) be issued and allows the operator to modify the parameters in the IGDSMSmm member. The system displays the current status of the parameters before it issues the WTOR. The modifications take effect for the current IPL, but the contents of the IGDSMSmm SYS1.PARMLIB member do not change. NO, which is the default, neither displays the parameters nor allows the operator to modify them.

If you specify both ID and PROMPT, enclose them in single quotation marks and separate them with a comma.

You have the option of manually starting SMS with a command or automatically starting it at future IPLs. When you are first preparing for SMS, omit the IGDSSIIN module name from the SMS record. This defines SMS to MVS at the next IPL but does not automatically start SMS. After you IPL the system, SMS is defined as a subsystem and you can then start it manually.

When you have completed preparations and are ready to start SMS, use the **T SMS=mm** command, where 'mm' identifies IGDSMSmm as the SMS initialization member. (The T SMS=mm command is an abbreviation for the SET SMS=mm command. To eliminate confusion with the SETSMS operator command, the abbreviated T SMS=mm form of the SET SMS=mm command is used throughout the remainder of this manual.)

When you have sufficiently tested your operations and are ready to have SMS automatically started at future IPLs, add the IGDSSIIN module name to the SMS record.

Examples of the SMS Record:

- An SMS record containing only:

  SMS

  defines SMS to MVS but does not start the subsystem at future IPLs. The ID and PROMPT values have no significance in this situation.

- An SMS record containing:

  SMS,IGDSSIIN

  defines SMS to MVS and starts the subsystem at future IPLs. The system uses the default values of ID and PROMPT: IGDSMS00 serves as a repository for initialization information, and the operator has no control over the rest of SMS initialization.

- An SMS record containing:

  SMS,IGDSSIIN,PROMPT=YES

  defines SMS to MVS and starts the subsystem at future IPLs. The system uses the default value of ID, which identifies IGDSMS00 as a repository for initialization information. The PROMPT parameter requests that a WTOR be issued and allows the operator to modify the parameters in IGDSMS00.

- An SMS record containing:

  SMS,IGDSSIIN,'ID=01,PROMPT=DISPLAY'

  defines SMS to MVS and starts the subsystem at future IPLs. The system uses IGDSMS01 as a repository for initialization information. The PROMPT parameter requests that the contents of IGDSMS01 be displayed, but the operator cannot modify them.

Accessing the Storage Administrator Primary Option Menu

The first time you select ISMF, you get the ISMF Primary Option Menu for *end users*. To get the ISMF Primary Option Menu for *storage administrators* (which is shown in Figure 2 on page 3), select option 0, ISMF PROFILE, and press ENTER. Within the ISMF Profile Option Menu, select option 0, USER MODE, and press ENTER. You get the User Mode Entry Panel, where you indicate that you want the storage administrator Primary Option Menu for all future ISMF sessions. To do this, select option 2 on the User Mode Entry Panel. After changing the user mode, you must exit ISMF and then return to it to view the Primary Option Menu for Storage Administrators.

Chapter 12, "Protecting SMS" on page 109 explains how to prevent end users from gaining access to the storage administrator Primary Option Menu through the ISMF PROFILE option.

Chapter 3. Defining SCDS Base Configuration Information

The systems in an SMS complex manage storage using the ACDS that you identify in IGDSMSmm. The ACDS contains a copy of the most recently activated SCDS. Each SCDS contains one SMS configuration, which defines a storage management policy. An SMS configuration consists of a base configuration, SMS class and storage group definitions, and ACS routines.

This chapter describes base configurations and explains how to define their contents using the Control Data Set Application of ISMF.

Understanding SCDSs

After allocating an SCDS as described in "Allocating Control Data Sets" on page 7, you can define its contents to contain a storage management policy. Before its contents can become the active storage management policy for an installation, an SCDS must satisfy certain validation requirements. The status of an SCDS can be either **valid** or **invalid**:

Valid   All ACS routine references to SMS classes and storage groups are consistent with the SMS class and storage groups definitions. Minimally, the SCDS must contain a:

- Storage class definition
- Pool storage group definition with at least one volume
- Storage group selection routine
- Base configuration.

If you specify a management class selection routine in an SCDS, then you must also specify management class definitions. If you specify a data class selection routine in an SCDS, then you must also specify data class definitions.

Invalid   At least one ACS routine reference to an SMS class or storage group is inconsistent with the SMS class or storage group definition, or one or more of the four required items for a valid SCDS are missing (storage class definition, pool storage group definition with at least one volume, storage group selection routine, and base configuration).

You can only activate an SMS configuration SCDS that is in a valid status. Activation validates the SCDS contents and copies the contents into the ACDS identified by IGDSMSmm. You cannot activate an SMS configuration SCDS that is in an invalid status.

"Validating ACS Routines or an Entire SCDS" on page 68 provides additional information on the validation of individual ACS routines and entire SCDSs.

Understanding Base Configuration Information

The first thing that you define in an SCDS is base configuration information. A base configuration contains installation defaults, such as a default management class, and identifies the systems to which the SMS configuration applies.

Planning SCDS Base Configuration Information

You need to determine the system names that you want to specify in the SCDS. These systems will constitute the SMS complex.

Before defining a base configuration, consider what you want to do with SMS-managed data sets that do not have a management class. You can specify a default management class for these data sets in the base configuration information (after defining the management class through the Management Class Application Selection option of ISMF, which is described in Chapter 5, "Defining Management Classes" on page 37), but a definition for the default management class must exist in the configuration.

The Starter Set in Appendix B, "Starter Set" on page 143, contains a NOACTION management class that is useful as a default management class. If the management class ACS routine does not determine a management class for a data set, DFHSM processes the data set using the default management class, if one exists.

You need to decide if SMS and non-SMS JCL are to be consistent when allocating data sets. If you decide to make them consistent, then you can define a default unit as part of the base configuration information, rather than requiring users to specify the UNIT parameter on their JCL.

See *MVS SML: Storage Management Subsystem Migration Planning Guide* for additional planning information.

Defining SCDS Base Configuration Information

After allocating an SCDS, you can define a base configuration. By selecting option 8 from the ISMF Primary Option Menu for Storage Administrators, you can invoke the Control Data Set Application Selection panel shown in Figure 6 on page 19.

```
                             CDS APPLICATION SELECTION
 COMMAND ===>

 TO PERFORM CONTROL DATA SET OPERATIONS, SPECIFY:

   CDS NAME ===> 'SMS.SCDS1.SCDS'
                               (1 to 44 character data set name)

 SELECT ONE OF THE FOLLOWING OPTIONS ===> 1

   1  DEFINE  - Define the Base Configuration
   2  ALTER   - Alter the Base Configuration










 USE ENTER TO PERFORM SELECTION;
 USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 6. Control Data Set Application Selection

In this panel, you can specify the name of the SCDS that is to contain the base
configuration. ISMF primes the CDS NAME field with the most recently refer-
enced SCDS name.

After specifying an SCDS, select option 1, DEFINE, and press ENTER to see the
SCDS Base Define panel shown in Figure 7. If you want to look at or alter an
SCDS that you defined previously, select option 2, ALTER, from the Control Data
Set Application Selection panel and press ENTER. The Alter panel contains the
same fields as the Define panel.

```
                             SCDS BASE DEFINE
 COMMAND ===>

 SCDS NAME:    SMS.SCDS1.SCDS
 SCDS STATUS:  INVALID

 TO ALTER SCDS BASE, SPECIFY:
   DESCRIPTION ===> THIS IS A CDS FOR OAM
               ===>
   DEFAULT MANAGEMENT CLASS ===>            (1 to 8 characters)
   DEFAULT UNIT             ===>            (esoteric or generic device name)
   DEFAULT DEVICE GEOMETRY
     BYTES/TRACK            ===> 47476      (1-999999)
     TRACKS/CYLINDER        ===> 15         (1-999999)

   SPECIFY ONE OF THE FOLLOWING OPTIONS ===>    (1 Add, 2 Delete, 3 Rename)
           SYSTEM NAME       ===>
           NEW SYSTEM NAME   ===>                (For option 3, Rename)

   SYSTEMS:  SYSTEM01
   SYSTEMS:

 USE ENTER TO PERFORM VERIFICATION;
 USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 7. SCDS Base Define

As its name implies, an SCDS base configuration forms the foundation for building an SCDS. The fields of this panel are described below:

SCDS NAME

This output field contains the SCDS name that you specified in the Control Data Set Application Selection panel.

SCDS STATUS

This output field tells you whether the SCDS is valid or invalid. Because you have not yet defined any ACS routines or SMS classes or storage groups, the SCDS STATUS reads INVALID. SMS sets the status when you save an SCDS with the END command and when you issue the VALIDATE command from the CDS Application Selection Panel.

Material in Figure 36 on page 71 lists conditions that cause an SCDS to be invalid, and "Validating ACS Routines or an Entire SCDS" on page 68 provides more information about the VALIDATE command.

DESCRIPTION

This is an optional free format input field of up to 120 characters. The default is blank.

DEFAULT MANAGEMENT CLASS

For SMS-managed data sets that have not been assigned a management class, DFHSM uses the default management class for expiration, migration, and backup information. However, the default management class name is not saved in the data set catalog entry. Specifying a default management class is optional. The default is blank.

DEFAULT UNIT

The default unit is an esoteric or generic device name, such as SYSDA or 3380, that applies to data sets that are not SMS-managed. For allocations, the default unit makes the JCL for these data sets consistent with the JCL for SMS-managed data sets.

If end users fail to specify the UNIT parameter on their JCL, expecting the data set will be SMS managed, SMS applies the default unit to all data sets having a disposition of either MOD or NEW that are not SMS managed. If you specify a default unit, make certain that it exists on the system performing the allocations. If you do not specify a default unit, end users must code the UNIT parameter to allocate data sets that are not SMS managed. The default is blank.

DEFAULT DEVICE GEOMETRY

The default device geometry isolates users from the actual physical device where SMS places their data sets. The amount of space will be consistent no matter what device is chosen to satisfy the allocation. SMS uses the following calculated space:

- As input to the Automatic Class Selection which determines the constructs
- For determining the number of tracks or cylinders on which to allocate the data set.

When allocating space for a new data set, SMS converts all requests for space into KB or MB using the actual geometry. For a generic device type such as

the 3380, SMS uses the actual device geometry for the allocation. For an esoteric device type such as SYSDA, SMS uses the default device geometry for the allocation. If end users fail to specify the UNIT parameter for SMS data sets, SMS uses the default device geometry for the allocation. If the end users in your installation use esoteric UNIT specifications, you must specify a default device geometry.

You will define the default device geometry when building the system configuration by specifying the number of bytes per track and the number of tracks per cylinder for converting to bytes.

BYTES/TRACK represents the number of bytes per track that SMS uses on allocations. The default is 47476 bytes for a 3380.

TRACKS/CYLINDER represents the number of tracks per cylinder that SMS uses on allocations. The default is 15 tracks for a 3380.

SYSTEMS IN THE COMPLEX

You can add, delete, or rename the systems defined to an SMS complex. SYSTEMS displays, in alphabetical order, the systems that you have previously defined. You must add at least one system name to a base configuration before you can save the SCDS, and can define up to eight systems to the SMS complex.

The system names must match the value of the SYSNAME parameter in the IEASYSyy member of the respective SYS1.PARMLIB. All fields have a default value of blank.

To add a system to the base configuration: Choose option 1, **Add**, and enter the system name in the SYSTEM NAME field. Press ENTER to add the system.

To delete a system from the base configuration: Choose option 2, **Delete**, and enter the system name in the SYSTEM NAME field. Press ENTER to get the Confirm Delete Request panel, where you can then delete the system name from the SCDS.

To rename a system in a base configuration: Choose option 3, **Rename**, enter the current name of the system in the SYSTEM NAME field, and enter the new name of the system in the NEW SYSTEM NAME field. Press ENTER to rename the system.

After defining base configuration parameters, you can verify their completeness and correctness by pressing ENTER. If your SCDS base configuration contains any errors, the cursor moves to the error and an error message appears in the short message area. Correct any errors and press ENTER to verify the new contents of the SCDS base configuration.

After correcting all errors, use the END command to save the SCDS base configuration and to leave the SCDS Base Define/Alter panel. A message indicates that SMS has saved your SCDS base configuration.

If at any time you want to leave either the SCDS Base Define or Alter panel without saving the changed base configuration information, issue the CANCEL command.
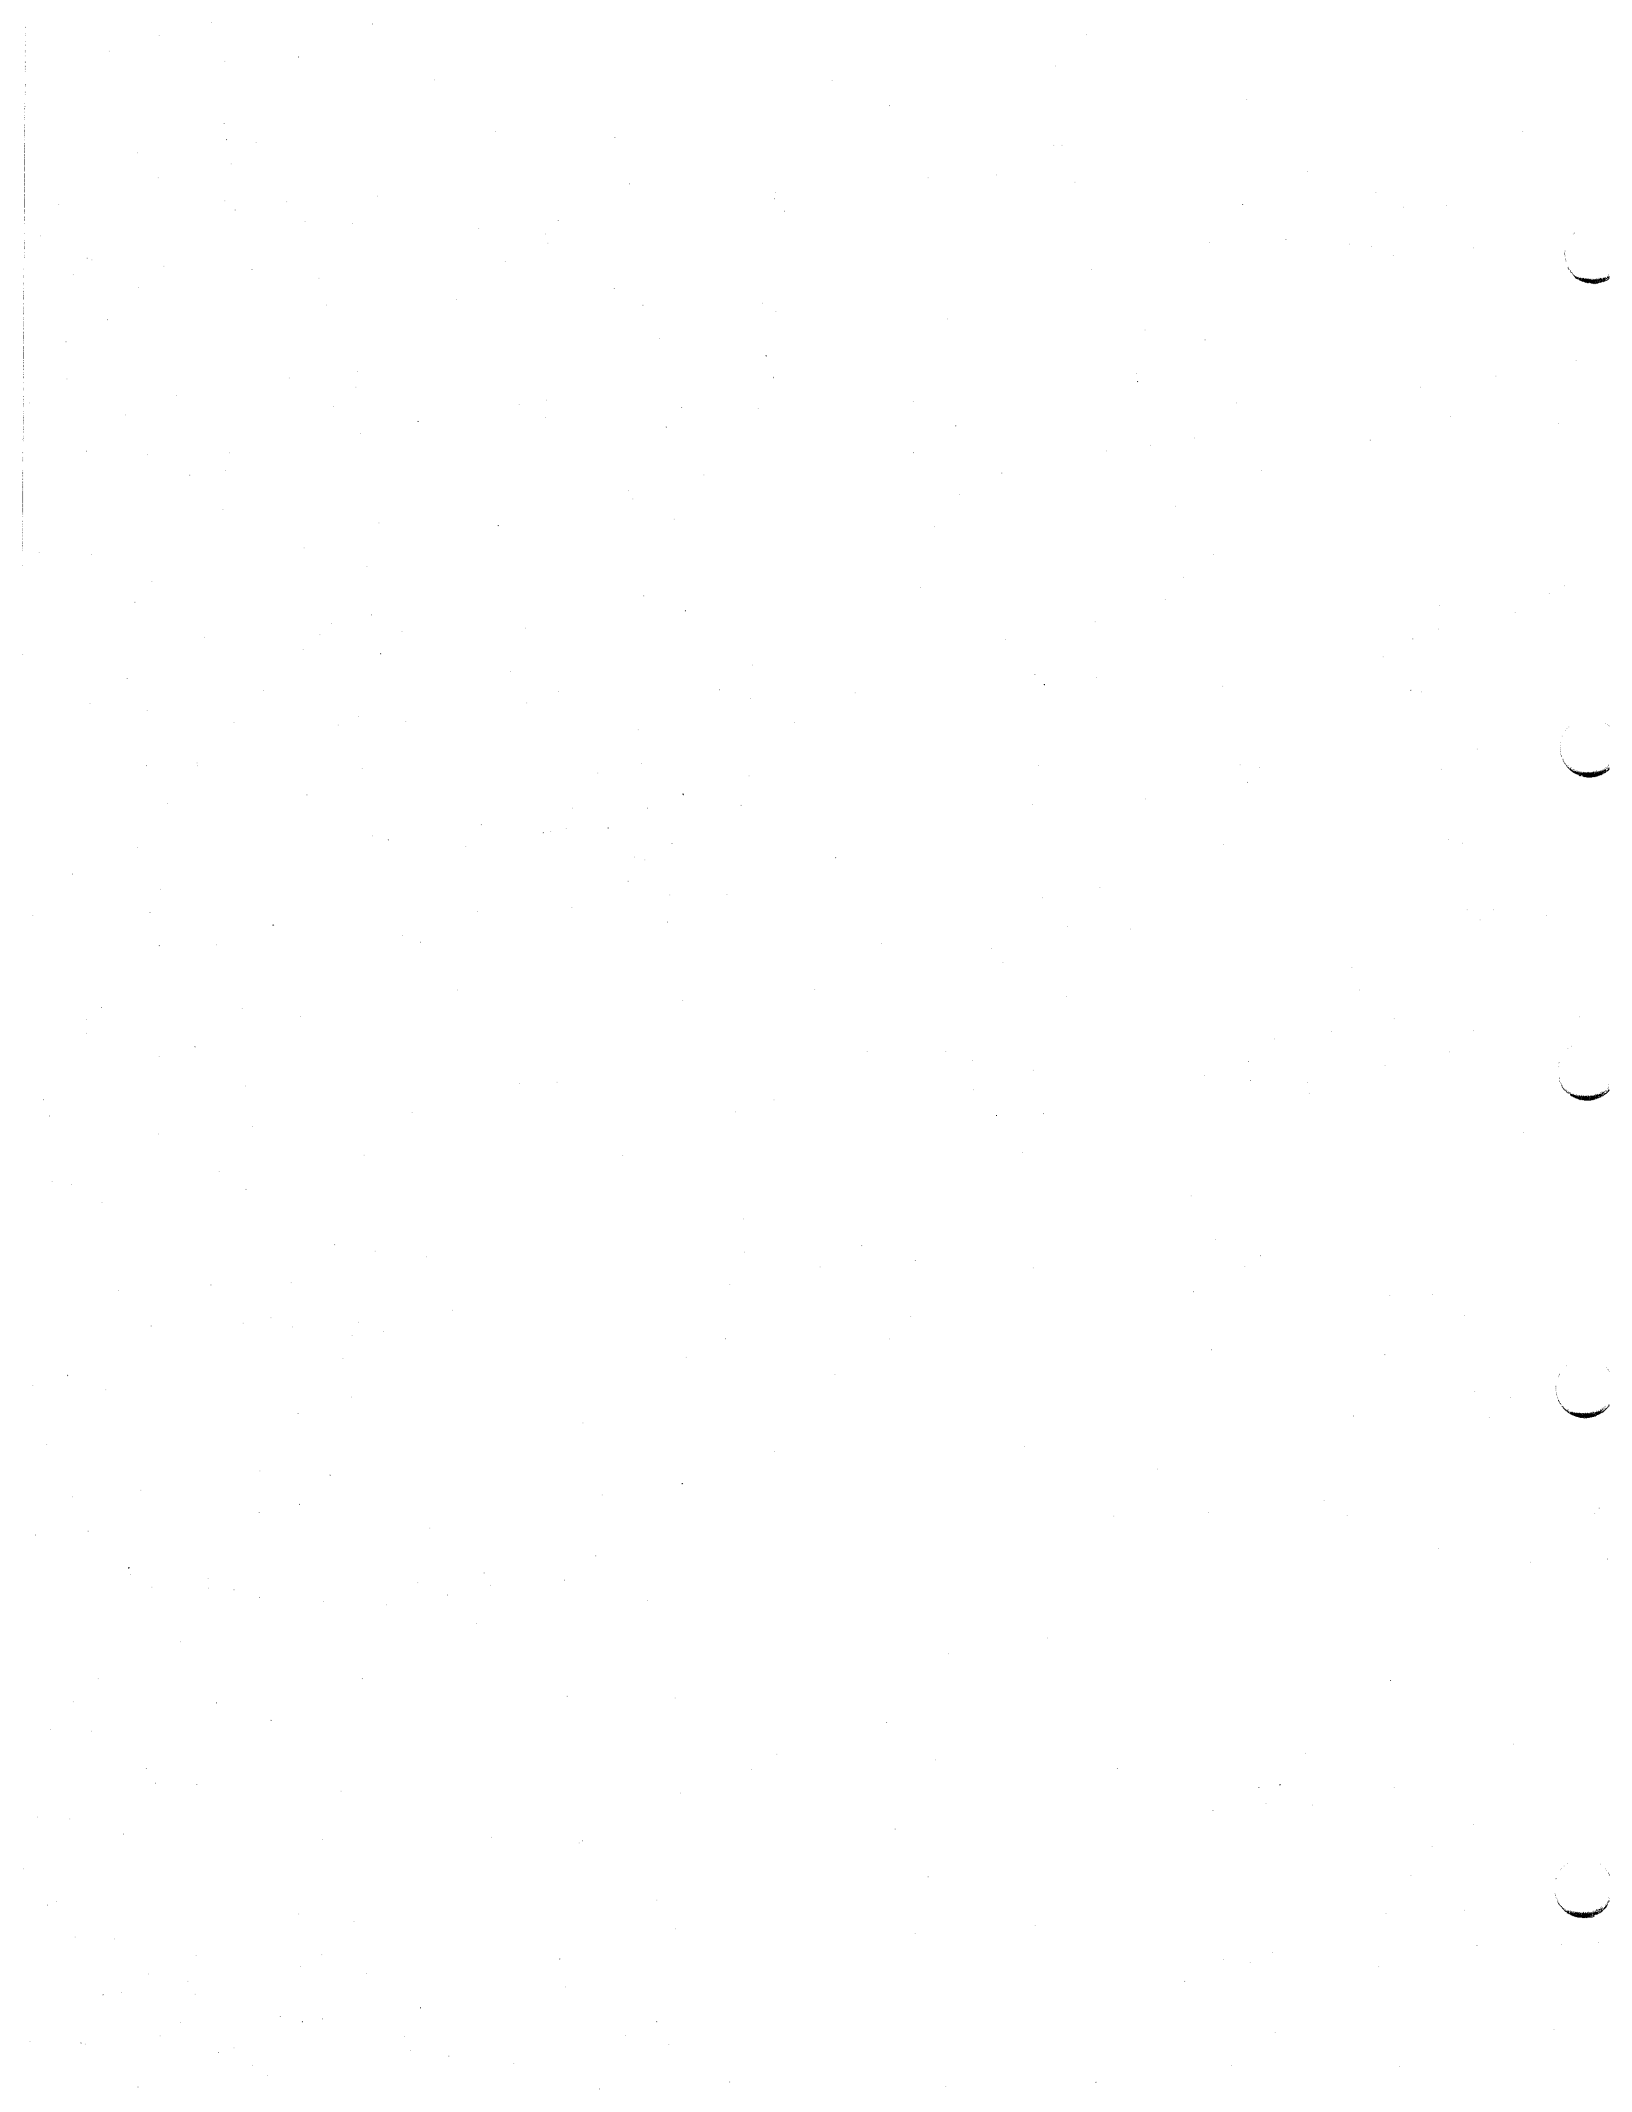
Chapter 4. Defining Storage Groups

When managing DASD volumes, you previously viewed and maintained them as individual devices. If you had too many critical data sets on a volume, you had to spread the data sets across other volumes to remove I/O bottlenecks. If you had too many low activity data sets on a volume, you might have been wasting one valuable resource and causing trouble for many others. In these and other instances, you had to work on individual volumes to solve your problems. SMS simplifies the management of DASD devices by pooling them together in storage groups. This chapter describes storage groups and shows you how to define them using the Storage Group Application of ISMF.

Understanding Storage Groups

A **storage group** can be either VIO or a list of DASD volumes. When an end user creates a data set, the storage group ACS routine builds a list of from one to fifteen storage groups. From the list of storage groups, SMS then selects an eligible volume for the data set (or volumes for multivolume data sets). Storage group definitions are transparent to end users. Only you as storage administrator can define, alter, or display their contents.

When defining storage groups, you must ensure that the actual, physical paths connecting systems to volumes match the desired, logical paths specified in the storage group definitions. Merely establishing a physical connection from a system to a volume does not provide access. Within the storage group definition, you must specify which systems have access to which storage groups, and which storage groups have access to which volumes. Likewise, merely defining a system to have access to a volume does not establish a physical connection. You need to ensure that the physical connection exists.

Two different storage groups cannot share a volume. You must define an entire volume to a storage group. You cannot have half of the volume belong to one storage group, and the other half belong to a different storage group. Also, a data set can only reside in one storage group. A data set can span volumes within a single storage group, but it cannot span volumes belonging to several different storage groups. For VSAM, the entire sphere (base cluster and all alternate indexes) must be in the same storage group.

A storage group can contain any DASD device type that MVS supports, except for 3330V volumes. All the volumes in a storage group must have the same device geometry: the same number of bytes per track, and tracks per cylinder.

A storage group refers to its volumes by using their serial numbers rather than their physical hardware addresses. Each volume in a storage group must contain a VVDS and an indexed VTOC. SMS automatically creates the VVDS after allocating the first SMS data set to a volume.

Tracking Physical Volume Status

To prepare new volumes for SMS storage groups, use the Device Support Facilities (ICKDSF) INIT command to assign new volume serial numbers and allocate an indexed VTOC. The STGR keyword converts a volume to make it available for allocation of new SMS data sets. The indicators show that the volume is converted and available for allocation of new SMS-managed data sets.

You can invoke this function from the Volume Application and the Storage Group Application of ISMF. If you go to the Volume Application directly from the Primary Option Menu, you initialize the volume as a non-SMS volume. If instead you work from a volume list generated within the Storage Group Application, you initialize the volume as an SMS volume.

As you bring volumes under the control of SMS, you need to keep track of those that are ready to be system managed. The physical volume status can be CONVERTED, INITIAL, or (non-SMS). To display this information, on the Storage Group Application Panel, do a LISTVOL.

CONVERTED

A status of CONVERTED indicates that the volume is fully available for SMS control. To be converted, all data sets on the volume must have an associated storage class, and all permanent data sets on the volume must be cataloged in an integrated catalog facility catalog. The volume must have VTOC INDEX.

INITIAL

A status of INITIAL indicates that an attempt to convert the volume has been made, but the volume contains data sets that fail to satisfy the requirements of a CONVERTED volume. In the INITIAL state, no new allocations can occur on the volume. You cannot extend data sets to additional volumes. The INITIAL status allows DFDSS to be able to process the volume without having to issue a RESERVE for the entire duration of the conversion process. Also, to be eligible for conversion, data sets and the volume must have an organization that is supported by SMS. For more information, refer to the *MVS SML: Storage Management Subsystem Migration Planning Guide*.

An INITIAL volume has an indexed VTOC and a VVDS. Not all data sets on the volume are under SMS control.

Planning Storage Groups

Ideally, you would have only one large storage group containing all of your data, and the system would manage everything. Realistically, you have to account for a variety of data sets, such as data bases, large data sets, and temporary data sets.

To begin your planning, define one primary storage group and identify all the data sets that do not fit into this general storage group definition. Keep in mind your storage class and management class requirements.

In order for SMS to satisfy new data set allocation requests associated with storage classes specifying AVAILABILITY = CONTINUOUS, the storage class candidate storage groups must identify a sufficient number of devices attached through storage controls that support the dual copy feature. The dual copy feature must be enabled.

For SMS to meet storage class performance objectives with BIAS=WRITE specified, and with millisecond response objectives that cannot be met by non-DASD fast write devices, the storage class candidate storage groups must identify a sufficient number of devices attached through storage controls that support the DASD fast write feature. The DASD fast write feature must be enabled.

For SMS to meet storage class performance objectives with BIAS=READ specified, and with millisecond response objectives that cannot be met by non-cache devices, the storage class candidate storage groups must identify a sufficient number of devices attached through storage controls that support cache. The cache feature must be enabled.

See Appendix C, "Choosing Restricted Line Operators and List Commands" on page 159 for additional information on dual copy, DASD fast write, and cache.

If a data set has a management class that specifies automatic dump, backup, or migration, then you need to identify the storage group that has the volume containing the data set. You need to specify through ISMF the system that is to perform the dump, backup, or migration only if that particular system should do the processing. Otherwise, do not specify a particular system to do the processing. It will be done automatically. If a data set has a storage class that requires specific hardware, then you need to ensure that the data set is assigned to a storage group with the proper hardware.

Before you actually define your storage groups, you should gather the following information:

- Existing I/O hardware configuration
- Projected hardware requirements
- Estimated general requirements of user groups
- Anticipated security requirements
- Backup and recovery requirements
- Data set sizes
- Required system access, based on both shared data requirements and the needs of user groups.

Also, you must define a base configuration, as described in Chapter 3, "Defining SCDS Base Configuration Information" on page 17, before you define a storage group.

See *MVS SML: Managing Storage Pools* for more information about planning and implementing storage groups.

Defining Storage Group Attributes

You can use ISMF to define your storage groups by selecting option 6, STORAGE GROUP, from the ISMF Primary Option Menu for Storage Administrators. Figure 8 on page 26 illustrates the Storage Group Application Selection panel.

```
                    STORAGE GROUP APPLICATION SELECTION
COMMAND ===>

TO PERFORM STORAGE GROUP OPERATIONS, SPECIFY:

  CDS NAME           ===> 'SMS.SCDS1.SCDS'
                              (1 to 44 character data set name or 'ACTIVE')

  STORAGE GROUP NAME ===> VIO              (for Storage Group List, fully or
                                            partially specified or * for all)
  STORAGE GROUP TYPE ===> VIO              (VIO, POOL or DUMMY)

SELECT ONE OF THE FOLLOWING OPTIONS ===> 2

   1  LIST    - Generate a list of Storage Groups
   2  DEFINE  - Define a Storage Group
   3  ALTER   - Alter a Storage Group
   4  VOLUME  - Display, Define, Alter or Delete Volume Information

IF OPTION 1 CHOSEN ABOVE,
   RESPECIFY SORT CRITERIA            ===> N   (Y or N)

USE ENTER TO PERFORM SELECTION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 8. Defining a Storage Group

To define a storage group, you must specify a CDS NAME, STORAGE GROUP
NAME, and STORAGE GROUP TYPE on the panel and select option 2, DEFINE.
The CDS NAME must be the name of an SCDS. ISMF primes the field with the
most recently referenced name. (The default is the quoted word 'ACTIVE',
which represents the currently active configuration, but you cannot define the
storage groups for the 'ACTIVE' configuration.)

In the STORAGE GROUP NAME field, you must specify the name of the storage
group that you are defining. ISMF primes the field with the most recently refer-
enced name. The default is an asterisk, *, which represents all storage groups
in the specified CDS NAME.

In the STORAGE GROUP TYPE field, you must specify one of the following three
types:

VIO

> **Virtual I/O** (VIO) storage groups allocate data sets to VIO, which sim-
> ulates the activity of a DASD volume. VIO storage groups do not
> contain any actual DASD volumes.

Dummy

> **Dummy** storage groups contain the serial numbers of volumes that
> no longer reside on the system but that you want to treat as SMS
> volumes. This allows existing JCL that explicitly references the
> volumes in VOL=SER statements to work. If end users specify a
> VOL=SER in their JCL, and that serial number is in a dummy
> storage group list, then SMS issues a catalog request to find the
> desired data set (rather than using the volume serial number).
>
> A dummy storage group should not contain the serial number of a
> volume that exists in the system. If the volume does exist in the
> system and the data set is an SMS data set, no JCL errors occur but
> the job fails during allocation. If the volume does exist in the system

and the data set is not SMS managed, then the resulting errors depend on the type of data set. For uncataloged data sets, either the data set cannot be found or the wrong data set with the same name has been found. For cataloged data sets, the job fails during allocation.

Pool

> **Pool** storage groups contain the volume serial numbers of SMS-managed volumes.

After specifying a STORAGE GROUP TYPE, select option 2 (DEFINE) and press ENTER to reach the appropriate Storage Group Define panel (VIO, dummy, or pool). The remainder of this chapter provides the panel sequences in the Storage Group Application and explains the attribute values for each type of storage group. You can leave any of the Storage Group Define panels at any time without saving the storage group by issuing the CANCEL command.

Defining a VIO Storage Group

Figure 9 shows the VIO Storage Group Define panel.

```
                           VIO STORAGE GROUP DEFINE
  COMMAND ===>

  SCDS NAME:         SMS.SCDS1.SCDS
  STORAGE GROUP NAME: VIO

  TO DEFINE STORAGE GROUP, SPECIFY:

  DESCRIPTION ===> VIO MAXSIZE will vary according to the needs
              ===> of your installation.
  VIO MAXSIZE ===> 1024              (8 to 2000000)
  VIO UNIT    ===> 3380              (DEVICE TYPE)

  SYSTEM                    SYSTEM
  NAME          STATUS      NAME          STATUS
  --------      ------      --------      ------
  SYSTEM01 ===> ENABLE                ===>
           ===>                       ===>
           ===>                       ===>
           ===>                       ===>



  USE ENTER TO PERFORM VERIFICATION;
  USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 9. Defining VIO Storage Group Attributes

SCDS NAME and STORAGE GROUP NAME are output fields that contain the SCDS and storage group names you specified in the Storage Group Application Selection panel. DESCRIPTION is an optional field of 120 characters where you can describe the VIO storage group.

To define a VIO storage group, you must specify a VIO MAXSIZE and VIO UNIT. You must also assign a storage group STATUS for each SYSTEM NAME.

VIO MAXSIZE

VIO MAXSIZE is a required value that represents the maximum size, in kilo-bytes, of data sets to be allocated to VIO. If a data set exceeds the maximum size, then the allocation fails. However, if the storage group ACS routine assigns both a VIO storage group and a pool storage group to the list of candidate storage groups for the data set, and the data set exceeds VIO MAXSIZE, then the data set is allocated to the pool storage group. ISMF primes the field with blanks, but you must specify a value.

VIO UNIT

VIO UNIT is a required field that represents the generic device type that this storage group simulates. ISMF primes the field with blanks, but you must specify a value.

SYSTEM NAME

SYSTEM NAME is an output field, which ISMF primes with the systems you specified in the SCDS Base Define and Alter panels.

STATUS is the current relationship between the VIO storage group and a given system. ISMF primes each field with ENABLE, which is the default.

After specifying the attribute values and system status relationships, issue the END command. END saves the newly defined VIO storage group and returns you to the Storage Group Application Selection panel where you will get the short message:

VIO saved

System Access to Storage Groups

If you want a system to access the volumes belonging to a storage group, you must ensure that the volumes are physically connected to the system. One of the following six relationships exists between each system in the SMS complex and a given VIO or pool storage group.

ENABLE   The system can create and access data sets on any volume belonging to the storage group. ENABLE is the default relationship between a system and a storage group.

DISALL   Disable All prevents the system from creating or accessing data sets on any volume belonging to the storage group.

DISNEW   Disable New prevents the system from creating new data sets (and DISP=MOD data sets that do not currently exist) on any volume belonging to the storage group.

QUIALL   Quiesce All prevents the system from scheduling jobs that create or access data sets on any volume belonging to the storage group. This state only affects systems that support job scheduling.

QUINEW   Quiesce New prevents the system from scheduling jobs that create new data sets and from modifying existing data sets (DISP=MOD) on any volume belonging to the storage group. This state only affects systems that support job scheduling.

NOTCON   Not Connected indicates that the storage group is defined but not accessible to the system. It resembles DISALL, except you cannot dynamically change the NOTCON system status. If you want a system group to have access to NOTCON volumes, physically connect them to the system, define the connection by changing the system status in the ISMF Storage Group Application, and activate the new configuration that defines the connection.

Defining a Dummy Storage Group

Figure 10 shows the Dummy Storage Group Define panel.

```
                           DUMMY STORAGE GROUP DEFINE
COMMAND ===>

SCDS NAME:            SMS.SCDS1.SCDS
STORAGE GROUP NAME:  DUMMY

TO DEFINE STORAGE GROUP, SPECIFY:

 DESCRIPTION ===>
             ===>













USE ENTER TO PERFORM VERIFICATION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 10. Defining Dummy Storage Group Attributes

SCDS NAME and STORAGE GROUP NAME are output fields containing the SCDS and storage group names that you specified in the Storage Group Application Selection panel. DESCRIPTION is an optional field of 120 characters where you can describe the dummy storage group. ISMF primes the field with blanks, which are the default. Issue the END command to save the newly defined dummy storage group and return to the Storage Group Application Selection panel.

At the moment, the storage group contains no volume serial numbers. To add volume serial numbers to the storage group, select option 4, VOLUME, from the Storage Group Application Selection panel shown in Figure 8 on page 26. ISMF primes the other fields on the panel with the most recently specified values. (You do not need to specify a STORAGE GROUP TYPE.) Then, press ENTER to define volume serial numbers to the storage group in the Storage Group Volume Selection panel shown in Figure 11 on page 30.

On this panel, you specify the volume serial numbers that you want to belong to the dummy storage group.

```
┌─────────────────────────────────────────────────────────────────────┐
│                      STORAGE GROUP VOLUME SELECTION                   │
│  COMMAND ===>                                                         │
│                                                                       │
│  CDS NAME:           SMS.SCDS1.SCDS                                   │
│  STORAGE GROUP NAME: DUMMY                                            │
│  STORAGE GROUP TYPE: DUMMY                                            │
│                                                                       │
│  SELECT ONE OF THE FOLLOWING OPTIONS   ===> 2                         │
│                                                                       │
│     1  DISPLAY      - Display SMS Volume Status (Pool only)           │
│     2  DEFINE       - Add a Volume to Volume Serial Number List       │
│     3  ALTER        - Alter Specific Volume and Statuses   (Pool only)│
│     4  DELETE       - Delete a Volume from Volume Serial Number List  │
│                                                                       │
│  SPECIFY VOLUME SERIAL NUMBER            ===> 123456                  │
│                                                                       │
│                                                                       │
│                                                                       │
│                                                                       │
│                                                                       │
│                                                                       │
│  USE ENTER TO PERFORM SELECTION;                                     │
│  USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.                 │
└─────────────────────────────────────────────────────────────────────┘
```

Figure 11. Defining Dummy Storage Group Volume Serial Numbers

CDS NAME, STORAGE GROUP NAME, and STORAGE GROUP TYPE are output
fields that reflect the values you have specified on previous panels. To add a
volume to the dummy storage group, select option 2, DEFINE, and specify the
volume serial number in the SPECIFY VOLUME SERIAL NUMBER field. The
value can be from one to six characters, but it must be a fully specified volume
serial number. ISMF primes the field with blanks, which is the default.

After specifying a volume serial number, press ENTER to add the volume to the
dummy storage group. Define any additional volume serial numbers, then use
END to return to the Storage Group Application Selection panel.

Defining A Pool Storage Group

Figure 12 on page 31 shows the Pool Storage Group Define panel.

```
                           POOL STORAGE GROUP DEFINE
     COMMAND ===>

     SCDS NAME:            SMS.SCDS1.SCDS
     STORAGE GROUP NAME: PRIMARY

     TO DEFINE STORAGE GROUP, SPECIFY:

      DESCRIPTION  ===> Contains most all data sets.
                   ===>
      AUTO MIGRATE ===> Y  (Y or N)    MIGRATE SYSTEM NAME ===> SYSTEM01
      AUTO BACKUP  ===> Y  (Y or N)    BACKUP SYSTEM NAME  ===>
      AUTO DUMP    ===> Y  (Y or N)    DUMP SYSTEM NAME-   ===> SYSTEM01

      DUMP CLASS   ===> SUNDAY         (1 to 8 characters)
      DUMP CLASS   ===>                DUMP CLASS     ===>
      DUMP CLASS   ===>                DUMP CLASS     ===>

      ALLOCATION/MIGRATION THRESHOLD: HIGH ===> 90  (1-99)    LOW ===> 80  (1-99)

      DEFINE SMS STORAGE GROUP STATUS     ===> N   (Y or N)

     USE ENTER TO PERFORM VERIFICATION AND SELECTION;
     USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 12. Defining Pool Storage Group Attributes

SCDS NAME and STORAGE GROUP NAME are output fields containing the
SCDS and storage group names that you specified in the Storage Group Appli-
cation Selection panel. DESCRIPTION is an optional field of 120 characters
where you can describe the pool storage group. ISMF primes the field with
blanks, which are the default.

AUTO MIGRATE

Note: The storage group and management class are interrelated. For
example, if you set AUTOMIGRATE or AUTOBACKUP to NO in storage group,
the storage group attributes will override the management class migration and
backup attributes.

In AUTO MIGRATE, you specify whether you want the volumes in this storage
group to be eligible for automatic space management processing. It is a
required field, which ISMF primes with the value Y, yes. Y specifies that
DFHSM is to perform daily space management processing and is a necessary
condition to perform interval migration processing.

AUTO BACKUP

In AUTO BACKUP, you specify whether you want the volumes in this storage
group to be eligible for automatic backup processing. It is a required field,
which ISMF primes with the value Y, yes. Y specifies that DFHSM is to perform
data availability management processing.

AUTO DUMP

In AUTO DUMP, you specify whether you want to automatically dump all the
volumes in this storage group. It is an optional field, which ISMF primes with
the value N, no. If you specify a value of Y, yes, DFHSM performs automatic
dump processing.

MIGRATE, BACKUP, and DUMP SYSTEM NAME

Before specifying a system in any of the SYSTEM NAME fields on this panel, you should define the system to SMS in the base configuration (using the CDS Application). See "SYSTEMS IN THE COMPLEX" on page 21 for more information.

DFHSM on each host gathers all the storage groups that you select for automatic space management (AUTO MIGRATE = = = > Y). If migration is not restricted to being processed of the specified system names, then DFHSM on a different host by the specification of the MIGRATE SYSTEM NAME, then DFHSM performs automatic space management for volumes belonging to those storage groups and the storage groups that you leave blank.

DFHSM ignores storage groups for which you specify a different system name and it does not process volumes that have already been processed. Do not specify a SYSTEM NAME unless processing of the storage group for the function must be performed only on that one host because that limits the capabilities of DFHSM to perform the request.

The same rules apply to BACKUP SYSTEM NAME for data availability management processing, and to DUMP SYSTEM NAME for automatic dump processing. All three SYSTEM NAME fields are optional and primed with blanks.

Dump Classes

In the DUMP CLASS fields, you can specify up to five unique dump class names. When DFHSM dumps volumes that belong to the storage group, it directs their contents to the dump classes. To use dump classes, you must define their names and parameters using DFHSM. Then on this panel, you can identify the dump class names to SMS. These are optional fields that are primed with blanks.

Migration and Allocation Thresholds

In ALLOCATION/MIGRATION THRESHOLD, you specify an upper and lower space limit for the volumes in a pool storage group. Although SMS attempts to stay below this value, the HIGH threshold value is used by the SMS volume selection function to prevent allocating space for data sets which would cause this value to be exceeded for a given volume in a storage group. In addition, this HIGH threshold value is used by DFHSM to determine whether the Interval Migration function should be performed on a volume in the storage group. The LOW threshold value is used by DFHSM during the DFHSM Interval Migration function as the threshold goal in reducing the amount of space occupied on a volume in the storage group. The LOW threshold value is also used by DFHSM during its Daily Space Management function as the threshold goal in reducing the amount of space occupied on a volume in the storage group. The LOW threshold value must be less than or equal to the HIGH threshold value.

Both numbers are percentages of the total space on the volume. If you specify Y, yes, for AUTO MIGRATE, then you must specify both a high and low threshold.

Defining Storage Group Status

In the DEFINE SMS STORAGE GROUP STATUS field, indicate whether you want to change the status of the pool storage group with respect to a given system in the SMS complex. Initially, all of the status fields are ENABLE. Specify Y, yes, to change a status. Accept the default N, no, to leave all the status fields as ENABLE.

When you have completed all your entries in the Pool Storage Group Define panel, press ENTER. If you specified Y for DEFINE SMS STORAGE GROUP STATUS, you get the SMS Storage Group Status Define panel shown in Figure 13, where you can define the status of the storage group to each system in the SMS complex.

```
                       SMS STORAGE GROUP STATUS DEFINE
COMMAND ===>

SCDS NAME:            SMS.SCDS1.SCDS
STORAGE GROUP NAME:   PRIMARY

TO DEFINE STORAGE GROUP SYSTEM STATUS, SPECIFY:

      SYSTEM          SMS SG
      NAME            STATUS
      --------        ------
      SYSTEM01  ===> ENABLE              ( Possible SMS SG Status
                ===>                       values for each System
                ===>                       NOTCON, ENABLE, DISALL,
                ===>                       DISNEW, QUIALL, or QUINEW )
                ===>
                ===>
                ===>
                ===>



USE ENTER TO PERFORM VERIFICATION; USE END COMMAND TO VIEW PREVIOUS PANEL;
USE HELP COMMAND FOR HELP; USE CANCEL COMMMAND TO EXIT.
```

Figure 13. Defining Storage Group System Status

SCDS NAME and STORAGE GROUP NAME are output fields that contain the information you specified on the Storage Group Application Selection panel. SYSTEM NAME is also an output field. It lists the systems that you defined in the base configuration (using the CDS Application). SMS SG STATUS lists the relationship between the storage group and each system in the SMS complex. One of the following six relationships exists between each system in the SMS complex and a given pool storage group:

- ENABLE
- DISALL
- DISNEW
- QUIALL
- QUINEW
- NOTCON

ENABLE is the default status. See "System Access to Storage Groups" on page 28 for an explanation of these relationships.

After establishing the relationships between your storage group and the systems in the SMS complex, you need to define volumes to the storage group.

From the SMS Storage Group Status Define panel, enter the END command to save your values and return to the Pool Storage Group Define panel. On this panel, Enter the END command to save the pool storage group.

This returns to you to the Storage Group Application Selection panel, which is primed with your CDS and storage group name. Select option 4, VOLUME, and press ENTER to view the Storage Group Volume Selection panel which appears in Figure 11 on page 30.

As they were for dummy storage groups, CDS NAME, STORAGE GROUP NAME, and STORAGE GROUP TYPE are output fields that contain the values you specified on the Storage Group Application Selection panel. Select option 2, DEFINE, specify a volume serial number, and press ENTER for each volume that you want to add to the pool storage group. Each time you press enter, you see the SMS Volume Status Define panel shown in Figure 14.

```
                              SMS VOLUME STATUS DEFINE
COMMAND ===>

SCDS NAME:              SMS.SCDS1.SCDS
STORAGE GROUP NAME:    PRIMARY
VOLUME SERIAL NUMBER:  123456

TO DEFINE SMS VOLUME STATUS, SPECIFY:

     SYSTEM        SMS VOL
     NAME          STATUS
     --------      -------
     SYSTEM01  ===> ENABLE           ( Possible SMS VOL Status
               ===>                    values for each System
               ===>                    NOTCON, ENABLE, DISALL,
               ===>                    DISNEW, QUIALL, or QUINEW )
               ===>
               ===>
               ===>
               ===>


USE ENTER TO PERFORM VERIFICATION; USE END COMMAND TO VIEW PREVIOUS PANEL;
USE HELP COMMAND FOR HELP; USE CANCEL COMMAND TO EXIT.
```

Figure 14. Defining Volume System Status

In the SMS Volume Status Define panel, you define the relationship between the volume in the pool storage group and each system in the SMS complex. SCDS NAME and STORAGE GROUP NAME are output fields that contain the information you specified on the Storage Group Application Selection panel. VOLUME SERIAL NUMBER is an output field that contains the serial number that you specified in the previous panel, Storage Group Volume Selection. SYSTEM NAME is also an output field that lists the systems you defined in the base configuration using the CDS Application.

Storage Group Access to Volumes

SMS VOL STATUS lists the relationship between a volume and each system in the SMS complex. Volumes have both a physical (actual) and a logical (system defined) connection to systems in the SMS complex. You are responsible for maintaining consistency between these two types of connections.

When a volume is physically connected to a system, you can define one of five different types of access (defined below) to that volume: ENABLE, DISALL, DISNEW, QUIALL, or QUINEW. When a volume is not physically connected to a system, you should define its logical connection to be NOTCON. The storage group containing the volume and all of the other volumes in the storage group must also be NOTCON with respect to the system. If you want a system to have access to a volume, you must ensure that the volume is both physically connected to the system and logically defined to it (through one of the five types of connected access).

One of the following six relationships exists between each system in the SMS complex and a given volume in the pool storage group:

ENABLE   The system can create and access data sets on the volume. ENABLE is the default relationship between a system and a volume.

DISALL   Disable All prevents the system from creating or accessing data sets on the volume.

DISNEW   Disable New prevents the system from creating new data sets (and DISP=MOD data sets that do not currently exist) on the volume.

QUIALL   Quiesce All prevents the system from scheduling jobs that create or access data sets on the volume. This state only affects systems that support job scheduling.

QUINEW   Quiesce New prevents the system from scheduling jobs that create new data sets or modify existing data sets on the volume. This state only affects systems that support job scheduling.

NOTCON   Not Connected indicates that the volume is defined but not accessible to the system. It resembles DISALL, except you cannot dynamically change the NOTCON system status. If you want a system to have access to the volume, you must ensure that the volume is physically connected to the system, define the connection in the ISMF Storage Group Application, and activate the new configuration that defines the connection.

When a system attempts to allocate a data set, it proceeds in order through the following checks with respect to the current system:

1. Storage group status—ENABLE/DISALL/DISNEW/QUIALL/QUINEW/NOTCON
2. SMS volume status—INITIAL/CONVERTED
3. MVS volume status—ONLINE/OFFLINE.

For a system to have access to a data set, you need to define one of the five types of access from the system to a storage group, and you need to define one of the five types of access from the storage group to a volume (ENABLE, DISALL, DISNEW, QUIALL, and QUINEW for both types of access). You must ensure that a physical connection exists from the volume to the system, and the MVS status of the volume must be ONLINE.

After defining the status of the volume, issue the END command to return to the Storage Group Application Selection panel. You can define additional volumes to the pool storage group, or you can issue the END command to return to the Storage Group Application Selection panel.

Assigning Storage Groups

Storage groups can only be determined through the storage group ACS routine. For a given SMS-managed data set, the storage group ACS routine is executed only if the storage class ACS routine assigned a valid storage class. The storage group ACS routine generates a list of one to fifteen storage groups, from which SMS selects an eligible volume for the data set (or volumes for multi-volume data sets). If the storage group ACS routine does not determine a storage group for a data set, allocation fails.

Defining Additional Storage Groups

You can copy existing storage groups and modify them to create new storage groups by using the COPY line operator, which is explained in "Copying SMS Classes and Storage Groups" on page 100. When you copy either a pool or VIO storage group within an SCDS, all of the system status values and system names for automatic processing are copied. When you copy either a pool or VIO storage group from one SCDS into a different SCDS, all of the system status values default to ENABLE rather than to the system status values of the source storage group. Also, the MIGRATE SYSTEM NAME, BACKUP SYSTEM NAME, and DUMP SYSTEM NAME are set to blank.

Chapter 5. Defining Management Classes

Formerly, DFHSM managed data sets at the volume level, applying the management criteria to all data sets on a given volume. SMS automates the management of storage at the data set level by introducing management classes. When assigned to data sets, management classes replace and expand attributes that are currently specified on JCL DD statements, IDCAMS DEFINE commands, and DFHSM commands. This chapter describes management classes and shows you how to define them using the Management Class Application of ISMF.

Understanding Management Classes

A management class is a list of data set migration, backup, and retention attribute values. DFHSM uses the attributes of the management class associated with a data set to manage storage. When you assign a management class to a data set, SMS places the management class name in both the basic catalog structure (BCS) and the VSAM volume data sets (VVDS) catalog entries of the data set. Management class is optional for SMS-managed data sets and does not apply to non-managed data sets.

If you alter a management class definition, SMS applies the changes to any new data sets after you activate the changed configuration. SMS also applies the changes to any previously created data sets, beginning with the next scheduled management cycle (such as daily space management or backup).

You can specify a default management class in an SCDS. DFHSM applies the default management class to all SMS-managed data sets that do not already have a management class. Unlike the data sets that you have already assigned a management class, the catalog entries of these data sets do not contain a management class name.

Describing Management Classes

A management class definition contains both descriptive and storage-related information. To identify and refer to your management classes, you assign each one a unique name that contains from one to eight alphanumeric characters. Each management class definition maintains an owner ID that identifies the storage administrator who originally created or last modified the management class. Also, each management class contains an optional 120 character description field for describing its contents.

Planning Management Classes

You should create a management class for each type of service that is to be provided by the installation. This type of service is a collection of data sets that have similar migration, backup, and retention requirements. Before you actually define your management classes, you should gather the following information:

- Release over allocation
- Migration requirements
- Retention criteria

- Treatment of expired data sets
- Frequency of backup
- Number of backup versions
- Retention of backup versions
- Generation data group (GDG) information.

Based on this information, you can establish management classes that centralize storage management in the SMS complex.

If a data set has a management class that specifies automatic backup or migration, then you must direct the data set to storage groups that are eligible to be processed for automatic backup or migration capabilities.

SMS manages storage at the data set level through management class, but a generation data group is managed differently. A generation group is a group of related cataloged data sets that have sequentially ordered names. The generation data groups are managed according to relative position in the GDG, such as latest to earliest. ACS filters a GDG rather than the base name. This prevents the user from specifying a different management class for a different generation data set in the same GDG base.

Defining Management Class Attributes

You can use ISMF to define your management classes by selecting option 3, MANAGEMENT CLASS, from the ISMF Primary Option Menu for Storage Administrators. Figure 15 illustrates the Management Class Application Selection panel.

```
                MANAGEMENT CLASS APPLICATION SELECTION
COMMAND ===>

TO PERFORM MANAGEMENT CLASS OPERATIONS, SPECIFY:

  CDS NAME              ===> 'SMS.SCDS1.SCDS'
                                  (1 to 44 character data set name or 'ACTIVE')

  MANAGEMENT CLASS NAME ===> INTERIM       (For Management Class List, fully or
                                           partially specified or * for all)

SELECT ONE OF THE FOLLOWING OPTIONS ===> 3

   1  LIST    - Generate a list of Management Classes
   2  DISPLAY - Display a Management Class
   3  DEFINE  - Define a Management Class
   4  ALTER   - Alter a Management Class

IF OPTION 1 CHOSEN ABOVE,
   RESPECIFY SORT CRITERIA           ===> N     (Y or N)


USE ENTER TO PERFORM SELECTION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 15. Defining a Management Class

To define a management class, you must specify a CDS NAME and a MANAGEMENT CLASS NAME on the panel and select option 3, DEFINE. The CDS NAME must be the name of an SCDS. ISMF primes the field with the most recently referenced name. (The default is the quoted word 'ACTIVE', which represents

the currently active configuration, but you cannot define the management classes to the 'ACTIVE' configuration.)

In the MANAGEMENT CLASS NAME field, you must specify the name of the management class that you are defining. Your management class names should be generic, for example, STANDARD, INTERIM rather than specific, for example, RETAIN30, BUDAILY. This is so that parameter changes do not make names meaningless or misleading. ISMF primes the field with the most recently referenced name.

Defining Management Class Expiration Attributes

Figure 16 shows the first page of the Management Class Define panel. You can leave any page of the Management Class Define panel without saving the management class by issuing the CANCEL command.

```
┌────────────────────────────────────────────────────────────────────────────┐
│                          MANAGEMENT CLASS DEFINE            Page 1 of 3      │
│  COMMAND ===>                                                                │
│                                                                             │
│  SCDS NAME:              SMS.SCDS1.SCDS                                      │
│  MANAGEMENT CLASS NAME:  INTERIM                                            │
│                                                                             │
│  TO DEFINE MANAGEMENT CLASS, SPECIFY:                                        │
│                                                                             │
│    DESCRIPTION ===>                                                          │
│                ===>                                                          │
│                                                                             │
│    EXPIRATION ATTRIBUTES                                                     │
│      EXPIRE AFTER DAYS NON-USAGE ===> 2        (1 to 9999 or NOLIMIT)       │
│      EXPIRE AFTER DATE/DAYS       ===> 7        (0 to 9999, yyyy/mm/dd       │
│                                                  or NOLIMIT)                 │
│                                                                             │
│    RETENTION LIMIT               ===> NOLIMIT   (0 to 9999 or NOLIMIT)      │
│                                                                             │
│                                                                             │
│                                                                             │
│  USE ENTER TO PERFORM VERIFICATION; USE DOWN COMMAND TO VIEW NEXT PANEL;    │
│  USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.│
└────────────────────────────────────────────────────────────────────────────┘
```

Figure 16. Defining Management Class Expiration Attributes

This page contains the management class expiration attributes. The expiration attributes are processed before the migration attributes that you specify on the second page of the Management Class Define panel. SCDS NAME and MANAGEMENT CLASS NAME are output fields that contain the SCDS and management class names you specified in the Management Class Application Selection panel. DESCRIPTION is an optional field of 120 characters where you can describe the management class.

EXPIRATION ATTRIBUTES and RETENTION LIMIT

You use EXPIRATION ATTRIBUTES to determine DFHSM action for data set expiration and deletion. DFHSM deletes expired data sets during automatic space management processing. The EXPIRATION ATTRIBUTES are required values that indicate when a data set becomes eligible for expiration. You can base expiration criteria on a specific date, on the number of days since the data set was created, or on the number of days since the data set was last referenced.

EXPIRE AFTER DAYS NON-USAGE specifies how much time must elapse since last access before a data set becomes eligible for expiration. EXPIRE AFTER DATE/DAYS specifies an absolute date or period for a data set to become eligible for expiration. The default for both fields is NOLIMIT.

RETENTION LIMIT is a required value that limits the use of retention period (RETPD) and expiration date (EXPDT) values that are explicitly specified in JCL or that are derived from data class definitions. If the value of a user-specified RETPD or EXPDT is within the limits specified by RETENTION LIMIT, it is saved for the data set.

The default RETENTION LIMIT is NOLIMIT. If you specify zero, a user-specified or data class derived EXPDT or RETPD is not saved. If end users specify values that exceed the maximum period, the RETENTION LIMIT overrides not only their values but also the EXPIRATION ATTRIBUTES values. RETENTION LIMIT is saved.

DFHSM first determines a data set becomes eligible for expiration if it satisfies both of the expiration attributes. If one of the EXPIRATION ATTRIBUTES is NOLIMIT, then the other attribute must be satisfied. If both of the EXPIRATION ATTRIBUTES are NOLIMIT, the data set never expires. DFHSM determines if a data set has expired based on the expiration date found in the catalog entry of the data set. If an expiration date is not found, DFHSM uses the management class EXPIRATION ATTRIBUTES.

Data sets having the INTERIM management class defined in Figure 16 on page 39 become eligible for expiration when both of the following criteria are met; at least seven days since creation and not referenced in the last two days.

Figure 17 shows several combinations of retention attributes. This example illustrates which values DFHSM uses for space management processing. The values in **boldface** are the ones that are to be used in each instance.

Figure 17. Comparing Retention Period Attributes

Attribute	Case 1	Case 2	Case 3	Case 4	Case 5
RETENTION LIMIT	0	**50**	100	NOLIMIT	100
EXPIRE AFTER DAYS NON-USAGE	**50**	50	50	50	50
EXPIRE AFTER DATE/DAYS	**100**	100	100	100	**100**
RETPD/EXPT	60	60	**60**	**60**	

In the first case, the RETENTION LIMIT is zero, so DFHSM honors the values of 50 and 100 which are the values specified for EXPIRE AFTER DAYS NON-USAGE and EXPIRE AFTER DATE/DAYS and ignores any user-specified or data class values. In the second case, the management class expiration values are ignored because RETPD and EXPDT values have been specified or derived and the RETENTION LIMIT is nonzero. However, because the RETENTION LIMIT is less than the user-specified or data class values, 50 is saved and DFHSM uses it to calculate the expiration date.

In the third case, the user specified or data class values fall within the RETENTION LIMIT. Consequently, DFHSM uses the values of 60 and 60. In the fourth case, the RETPD and EXPDT are used because RETENTION PERIOD is NOLIMIT.

In the final case, because no user specified or data class derived values are available, DFHSM uses the values specified in the management class expiration attributes which are shown in Figure 18.

Defining Management Class Migration Attributes

To DFHSM, a data set occupies one of two distinct states in storage:

Primary Also known as **level 0**, the primary state indicates that end users can directly access a data set residing on an SMS-managed volume.

Migrated End users cannot directly access data sets that have migrated from the primary state to a migrated state. To be accessed, the data sets must be recalled to primary storage. A migrated data set can reside on either **migration level 1** (normally permanently mounted DASD) or **migration level 2** (normally tape).

A data set can move back and forth between these two states, and it can move from level 0 to migration level 2 (and back) without passing through migration level 1.

See *DFHSM: General Information* for background information on data set migration.

Figure 18 shows the second page of the Management Class Define panel, which contains the management class migration and GDG management attributes. SCDS NAME and MANAGEMENT CLASS NAME are output fields that contain the SCDS and management class names you specified in the Management Class Application Selection panel.

```
                          MANAGEMENT CLASS DEFINE            Page 2 of 3
  COMMAND ===>

  SCDS NAME:              SMS.SCDS1.SCDS
  MANAGEMENT CLASS NAME:  INTERIM

  TO DEFINE MANAGEMENT CLASS, SPECIFY:

    PARTIAL RELEASE           ===> Y            (Y or N)

    MIGRATION ATTRIBUTES
      PRIMARY DAYS NON-USAGE    ===>            (0 to 9999 or blank)
      LEVEL 1 DAYS NON-USAGE    ===>            (0 to 9999, NOLIMIT or blank)
      COMMAND OR AUTO MIGRATE   ===> NONE       (BOTH, COMMAND or NONE)


    GDG MANAGEMENT ATTRIBUTES
      # GDG ELEMENTS ON PRIMARY ===> 1          (0 to 255 or blank)
      ROLLED-OFF GDS ACTION     ===> EXPIRE     (MIGRATE, EXPIRE or blank)



  USE ENTER TO PERFORM VERIFICATION; USE UP/DOWN COMMAND TO VIEW OTHER PANELS;
  USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 18. Defining Management Class Migration Attributes

To define a management class, you must specify values for PARTIAL RELEASE and COMMAND OR AUTO MIGRATE. If you specify either BOTH or COMMAND for COMMAND OR AUTO MIGRATE, then PRIMARY DAYS NON-USAGE and LEVEL 1 DAYS NON-USAGE become required fields. PRIMARY DAYS

NON-USAGE has a default value of 2, and LEVEL 1 DAYS NON-USAGE has a default of 60. Also, if you specify NONE, the EXPIRATION attributes still apply. The other fields on this panel are optional.

In PARTIAL RELEASE, you indicate if you want the allocated but unused space of a data set to be eligible for automatic release. It applies to non-VSAM data sets. ISMF primes the field with the last used value, and the default is N, no. The PARTIAL RELEASE attribute is not checked during the DFHSM Interval Migration Function.

MIGRATION ATTRIBUTES

PRIMARY DAYS NON-USAGE represents the minimum number of days that must elapse since last access before a data set is eligible for migration. The default is 2. A nonzero number for LEVEL 1 DAYS indicates the number of days that must elapse before a data set can migrate from level 1 to level 2. The default is 60. COMMAND OR AUTO MIGRATE allows you to specify if a data set is eligible to be migrated by BOTH command and automatic processing, by COMMAND alone, or not at all (NONE).

If you do not want to migrate data sets that belong to a particular management class, specify NONE in the COMMAND OR AUTO MIGRATE field. The data sets remain on primary storage until they expire. If you want data sets to be eligible for migration directly from level 0 storage to level 2 tape, specify 0 in the LEVEL 1 DAYS NON-USAGE field. Otherwise, the data sets must first migrate to level 1 (the days spent on level 0 count toward the eligibility of moving to level 2).

GDG MANAGEMENT ATTRIBUTES

The GDG MANAGEMENT ATTRIBUTES indicate how many generations of a GDG to keep on primary storage and what to do with rolled-off generation data sets (GDS). The PRIMARY DAYS NON-USAGE attribute applies to all GDG elements, but elements that exceed the value you specify in # GDG ELEMENTS ON PRIMARY become eligible for early migration.

GDG ELEMENTS ON PRIMARY indicates how many of the most recent generations of a GDG to keep on primary storage. Generations that fall over the limit become eligible for early migration during the next run of daily space management

ROLLED-OFF GDS ACTION indicates whether to expire rolled-off GDS or to make them eligible for migration. If you specify MIGRATE, management class EXPIRATION attributes will be applied to the data set to determine if the data set should be deleted, and if not, the data set is eligible for migration. If you specify EXPIRE, the rolled-off GDS will be deleted unless this expiration value has not expired. If the GDS is not deleted, the GDS is determined to be eligible for migration based on the other management class migration attributes.

Both the # GDG ELEMENTS ON PRIMARY and the ROLLED-OFF GDS ACTION attributes are optional and have default values of blank. If left blank, no special treatment results from the data set being a GDS or a rolled-off GDS. That is, other management class attributes will be used to process the data set.

During daily space management, GDS that have been determined as eligible for migration are migrated in the following priority order:

1. GDS that are rolled-off
2. GDS that exceed the limit specified in # GDG ELEMENTS ON PRIMARY
3. Other GDS's are handled by the non-GDG algorithm which bases priority on a function of size and the length of time it has been eligible to migrate in PRIMARY DAYS NON-USAGE.

After specifying the migration attributes, issue the DOWN command to specify the backup attributes on the last page of the Management Class Define panel, which is shown in Figure 19.

Defining Management Class Backup Attributes

The third page of the Management Class Define panel contains the management class backup attributes. ADMIN OR USER COMMAND BACKUP is used to indicate whether data sets managed by this management class may be processed by DFHSM backup and if so, what categories of users can issue data set commands for them. A value of NONE means that neither command or autoprocessing is permitted. ADMIN means only storage administrators can issue commands against these data sets. BOTH means both storage administrators and users can issue commands against these data sets. If the value is other than NONE, then AUTO BACKUP specifies whether these data sets are eligible for automatic backup.

```
                       MANAGEMENT CLASS DEFINE              Page 3 of 3
 COMMAND ===>
 SCDS NAME:            SMS.SCDS1.SCDS

 MANAGEMENT CLASS NAME:   INTERIM

 TO DEFINE MANAGEMENT CLASS, SPECIFY:

   BACKUP ATTRIBUTES
     BACKUP FREQUENCY                  ===> 1      (0 to 9999 or blank)
     NUMBER OF BACKUP VERSIONS         ===> 2      (1 to 13 or blank)
       (DATA SET EXISTS)
     NUMBER OF BACKUP VERSIONS         ===> 1      (0 to 13 or blank)
       (DATA SET DELETED)
     RETAIN DAYS ONLY BACKUP VERSION   ===> 60     (1 to 9999, NOLIMIT
       (DATA SET DELETED)                           or blank)
     RETAIN DAYS EXTRA BACKUP VERSIONS ===> 30     (1 to 9999, NOLIMIT
                                                    or blank)
     ADMIN OR USER COMMAND BACKUP      ===> NONE   (BOTH, ADMIN or NONE)
     AUTO BACKUP                       ===> Y      (Y or N)


 USE ENTER TO PERFORM VERIFICATION; USE UP COMMAND TO VIEW PREVIOUS PANEL;
 USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 19. Defining Management Class Backup Attributes

The ADMIN OR USER COMMAND BACKUP and AUTO BACKUP fields are required, and they have default values of BOTH and Y, respectively. ADMIN OR USER COMMAND BACKUP indicates if both the end user and the storage administrator can issue command backups, if only the storage administrator can, or if neither of them can. AUTO BACKUP specifies if a data set is eligible for automatic backup processing.

If you specify Y in the AUTO BACKUP field or either BOTH or ADMIN in the
ADMIN OR USER COMMAND BACKUP field, the remaining fields on this panel
are required. Otherwise, the remaining fields are optional.

BACKUP FREQUENCY

BACKUP FREQUENCY specifies how many days must elapse before DFHSM can
back up data sets that have changed since the last backup. If you want to back
up changed data sets every time that DFHSM processes the volumes containing
them, specify 0. The default is 1.

NUMBER OF BACKUP VERSIONS

The NUMBER OF BACKUP VERSIONS fields specify the maximum number of
backup versions to retain for a data set. The default is 2 if the data set still
exists and 1 if it has been deleted.

RETAIN DAYS ONLY BACKUP VERSIONS

RETAIN DAYS ONLY BACKUP VERSION(DATA SET DELETED) indicates how
many days to keep the most recent backup version of a deleted data set,
starting from the day DFHSM detects it has been deleted. This attribute applies
only when a data set no longer exists on primary (level 0) or migrated (levels 1
and 2) storage. The default is 60.

RETAIN DAYS EXTRA BACKUP VERSIONS

RETAIN DAYS EXTRA BACKUP VERSIONS indicates how many days to keep
backup versions other than the most recent one, starting from the day backups
were created. It only applies when more than one backup version exists.
When a new backup version is created and the number of backup versions
already equals the value specified in the pertinent (existing or deleted)
NUMBER OF BACKUPS field, the oldest version is deleted. This attribute
applies whether the data set has been deleted or not. The default is 30.

Assigning Management Classes

You can assign management classes either through the management class
ACS routine or by explicit specification. For permanent SMS-managed data
sets, the management class ACS routine is executed only if the storage class or
other mechanism ACS routine determined a valid storage class. The manage-
ment class ACS routine is not executed for temporary data sets.

You can explicitly specify a management class on the following, but a manage-
ment class determined by the management class ACS routine takes preced-
ence:

- JCL DD statements
- TSO/E ALLOCATE command
- DFDSS COPY and RESTORE commands
- Access method services ALLOCATE, DEFINE, and IMPORT, commands
- Dynamic allocation requests
- ISPF/PDF data set allocation panels.

End users cannot override any of the attribute values that you assign to a data set through the management class ACS routine except for the expiration attributes. Note that this is on the condition that the retention limit is not exceeded or set to NOLIMIT. See "Defining Management Class Expiration Attributes" on page 39 for information about the order of precedence for expiration attributes.

The format for specifying a management class on a JCL statement is:

MGMTCLAS=management-class-name

The format for specifying a management class on a TSO/E command is:

MGMTCLAS(management-class-name)

The format for specifying a management class on an access method services command is:

MGMTCLAS(-) or MANAGEMENTCLASS(management-class-name)

For information on determining management classes through ACS routines, see Chapter 8, "Defining ACS Routines" on page 63.

Defining Additional Management Classes

You can copy existing management classes and modify them to create new management classes by using the COPY line operator, which is explained in "Copying SMS Classes and Storage Groups" on page 100.
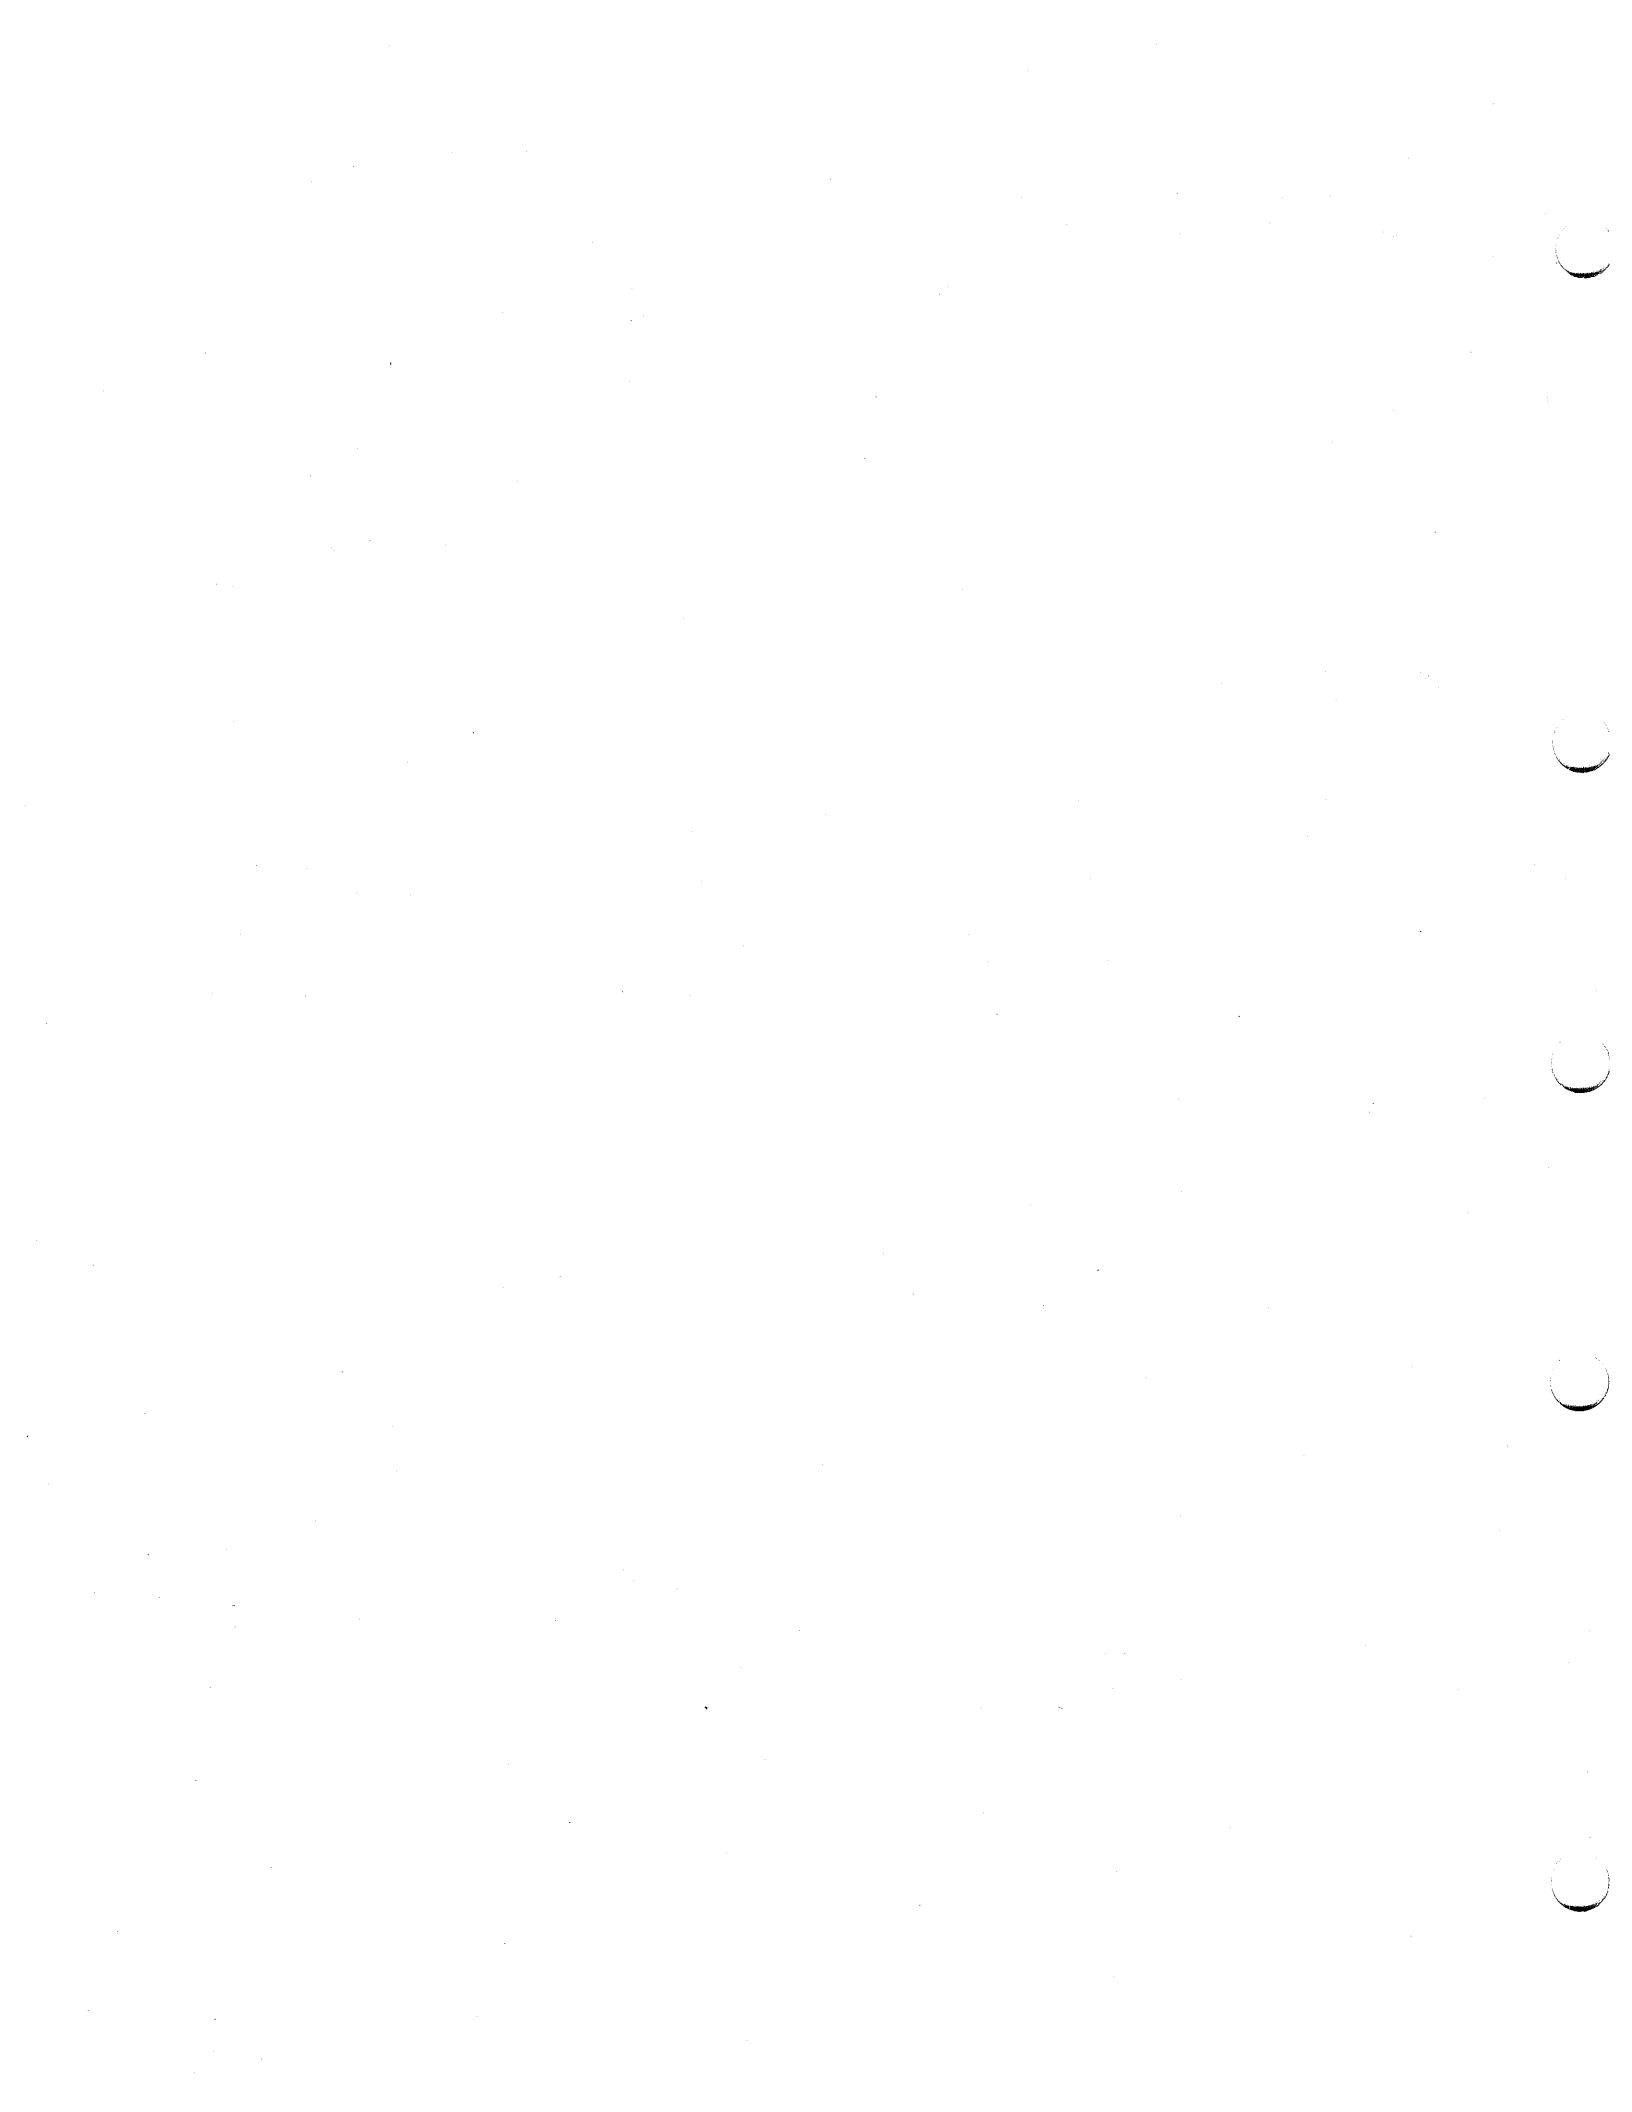
Chapter 6. Defining Storage Classes

When attempting to improve performance you previously placed important and critical data sets on selected storage devices. If you had data sets that consistently required short response times, you placed them on DASD volumes that had low I/O rates or that were connected to cache storage controls. If you had data sets that required continuous availability, you placed duplicate copies of them on other DASD volumes. SMS separates data set performance objectives and availability from physical storage by introducing storage classes. This chapter describes storage classes and shows you how to define them using the Storage Class Application of ISMF.

Understanding Storage Classes

A storage class is a list of performance objectives and availability requirements. Each storage class represents a list of services that are available to data sets having similar access requirements. A storage class does not represent any physical storage, but rather, provides the criteria that SMS uses in determining an appropriate volume on which to place a data set. With regard to performance objectives, SMS attempts to select a volume that meets or exceeds the specified objective, but SMS does not guarantee response time. If no volume satisfies the objective, then SMS makes one attempt to find a volume that most closely matches the specified objective. If no volume comes close enough to matching the specified objective, SMS selects the volume with the most available space.

To satisfy the availability requirements of critical data sets, SMS selects a volume that has the **dual copy** feature of the IBM 3990 Storage Control Model 3 active. With the dual copy feature active, the IBM 3990 Storage Control Model 3 maintains a duplicate copy of the selected volume. To satisfy the availability requirement of normal data sets, SMS selects a volume that does not have the dual copy feature active. Through storage class, you can also specify whether a data set is referenced primarily in the read or write mode. A storage class does not represent any specific physical storage.

Every data set that SMS manages has a storage class. When you assign a storage class to a data set, SMS places the storage class name in both the BCS and the VVDS catalog entries of the data set. Only you, as a storage administrator, can define or alter the storage classes for an SMS complex. End users cannot override the storage classes selected for a data set by the ACS routines.

See Appendix C, "Choosing Restricted Line Operators and List Commands" on page 159 for more information on establishing dual copy volumes and enabling or disabling cache functions. In this section, you can also find the IBM storage control units that support these features.

Describing Storage Classes

A storage class definition contains both descriptive and access-related information. To identify and refer to storage classes, you assign each one a unique name that contains from one to eight alphanumeric characters. Each storage class definition maintains an owner ID that identifies the storage administrator who originally created or last modified the storage class. Also, each storage class contains an optional 120 character description field for describing its contents.

Planning Storage Classes

You should create a storage class for each level of service that is to be provided by the installation. Before you actually define storage classes, you need to identify the available hardware in your installation. If you have data sets that require short response times, you need to have an IBM 3880 Storage Control Model 23 or an IBM 3990 Storage Control Model 3.

Defining Storage Class Attributes

You can use ISMF to define storage classes by selecting option 5, STORAGE CLASS, from the ISMF Primary Option Menu for Storage Administrators. Figure 20 illustrates the Storage Class Application Selection panel.

```
                      STORAGE CLASS APPLICATION SELECTION
COMMAND ===>

TO PERFORM STORAGE CLASS OPERATIONS, SPECIFY:

  CDS NAME            ===> 'SMS.SCDS1.SCDS'
                                    (1 to 44 character data set name or 'ACTIVE')

  STORAGE CLASS NAME ===> DBFAST        (For Storage Class List, fully or
                                         partially specified or * for all)


SELECT ONE OF THE FOLLOWING OPTIONS  ===> 3

  1   LIST     - Generate a list of Storage Classes
  2   DISPLAY - Display a Storage Class
  3   DEFINE   - Define a Storage Class
  4   ALTER    - Alter a Storage Class

IF OPTION 1 CHOSEN ABOVE,
   RESPECIFY SORT CRITERIA           ===> N  (Y or N)


USE ENTER TO PERFORM SELECTION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 20. Defining a Storage Class

To define a storage class, you must specify a CDS NAME and a STORAGE CLASS NAME on the panel and select option 3, DEFINE. The CDS NAME must be the name of an SCDS. ISMF primes the field with the most recently referenced name. (The default is the quoted word 'ACTIVE', which represents the currently active configuration, but you cannot define the storage classes to the 'ACTIVE' configuration.)

In STORAGE CLASS NAME, you must specify the name of the storage class that you are defining. ISMF primes the field with the most recently referenced name.

Defining Storage Class Performance and Availability Attributes

Select option 3, DEFINE, and press ENTER to reach the Storage Class Define panel shown in Figure 21. You can leave the Storage Class Define panel at any time without saving the storage class by issuing the CANCEL command.

```
                           STORAGE CLASS DEFINE
 COMMAND ===>

 SCDS NAME:            SMS.SCDS1.SCDS
 STORAGE CLASS NAME: DBFAST

 TO DEFINE STORAGE CLASS, SPECIFY:

   DESCRIPTION ===> This data base storage class provides the fastest
               ===> response time of all the data base storage classes.

   PERFORMANCE OBJECTIVES

     DIRECT MILLISECOND RESPONSE      ===> 5         (1 to 999 or blank)
     DIRECT BIAS                      ===> R         (R, W or blank)

     SEQUENTIAL MILLISECOND RESPONSE ===> 5          (1 to 999 or blank)
     SEQUENTIAL BIAS                  ===> R         (R, W or blank)

     AVAILABILITY                     ===> STANDARD   (STANDARD or CONTINUOUS)
     GUARANTEED SPACE                 ===> Y          (Y or N)

 USE ENTER TO PERFORM VERIFICATION;
 USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 21. Defining Storage Class Attributes

SCDS NAME and STORAGE CLASS NAME are output fields that contain the SCDS and storage class names you specified in the Storage Class Application Selection panel. DESCRIPTION is an optional field of 120 characters where you can describe the storage class.

To define a storage class, you must specify values in the AVAILABILITY and GUARANTEED SPACE fields. The remaining fields are optional.

In the PERFORMANCE OBJECTIVES, you can request MILLISECOND RESPONSE and indicate the bias of both direct and sequential access data sets.

MILLISECOND RESPONSE

The MILLISECOND RESPONSE represents an objective. SMS searches for a volume that potentially meets or exceeds this objective. If no volume satisfies the objective, then SMS makes one attempt to find a volume that most closely matches the specified objective. If no volume comes close enough to matching the specified objective, SMS selects the volume with the most available space. SMS uses this objective as a guideline but does not guarantee it.

BIAS

BIAS specifies if the majority of I/O is expected to be READ, WRITE, or unknown. The fields have a default of blanks, which represent an unknown bias. If you specify a read bias, R, you need to have cache storage control available. If you specify a write bias, W, you need to have the DASD fast write feature of the 3990 Storage Control Model 3 available. If you specify that data sets require caching (either an R or a W in BIAS), the data sets are cached if a control unit is available to the volume. If you specify a blank value for BIAS (unknown bias or equal balance between reads and writes), the data sets might still get cached if the response time is not 999. You can inhibit this caching by specifying a MILLISECOND RESPONSE of 999.

Note: Allocation is not failed if a caching or DASD fast write storage control is not available.

AVAILABILITY

The AVAILABILITY field is used to specify whether data set processing should continue after device failures. Specify an availability of CONTINUOUS if you want to continue to process a data set if a device failure severs communications with the volume that contains the data set. Specify an availability of STANDARD to allow processing of a data set to stop after device failures. STANDARD is the default.

You can use the dual copy feature of the 3990 Storage Control Model 3 to guarantee the availability of critical data sets, by specifying CONTINUOUS availability. The storage control maintains a duplicate copy of data written to a duplicate DASD volume. SMS only selects volumes that are backed by a dual copy. If none are available, allocation fails.

If data sets do not require such a high level of availability, specify STANDARD availability, which represents normal storage needs. SMS only selects volumes that are *not* dual copy. If none are available, allocation fails, even if some dual copy volumes have available space. STANDARD is the default.

GUARANTEED SPACE

You can use the GUARANTEED SPACE attribute to permit preallocation of space for multivolume data sets. Users can preallocate space for multivolume data sets by specifying a storage class with a guaranteed space attribute. SMS fails the request if insufficient space is available.

For a multivolume non-VSAM data set, space is preallocated on all the volumes. For a multivolume VSAM data set, space is allocated on the first volume. The remaining volumes, if explicitly specified, become candidate volumes to be used when the data set is extended. For VSAM key-sequenced data sets with key ranges specified, space is preallocated on as many volumes as key ranges specified.

The following JCL allocates a multivolume data set selected by SMS. Storage class DBLOG must have GUARANTEED SPACE. This example allocates 100Mb on each of five volumes. When all of the primary space is used, the 25Mb secondary space is allocated on the last volume.

```
//DD1  DD  DSN=ENG.MULTFILE,DISP=(,KEEP),STORCLAS=DBLOG
          SPACE=(1,(100,25)),AVGREC=M,
          UNIT=(3380,5)
```

In addition to preallocating space for multivolume data sets, you can also use the GUARANTEED SPACE attribute of storage class to allow a single or multiple volumes, explicitly specified by the user, to be honored by SMS. For example, you can allow users to preallocate a non-VSAM multivolume data set on specific volumes. You can also use the storage class GUARANTEED SPACE attribute to allocate space specifically on volumes within the same storage group. The following JCL preallocates space for a multivolume data set on volumes VOL001, VOL002, VOL003, VOL004, and VOL005 if the storage class SPECIAL has GUARANTEED SPACE:

```
//DD1   DD   DSN=ABCD.MULTFILE,DISP=(,KEEP),STORCLAS=SPECIAL
             SPACE=(1,(100,75)),AVGREC=M,
             VOL=SER(VOL001,VOL002,VOL003,VOL004,VOL005)
```

This JCL allocates 100Mb on all specified volumes for this data set. If all the space on these volumes is used by the data set, a secondary extent of 75Mb is allocated on the last volume, VOL005.

If you specify the default, N, no, for GUARANTEED SPACE, then SMS chooses the volumes for allocation, ignoring any VOL=SER statements specified on JCL. No space is preallocated.

End users can allocate space on specific volumes for a data set if you:

- Create at least one storage class with the the GUARANTEED SPACE attribute
- Ensure the user is authorized to the storage class with the GUARANTEED SPACE attribute
- Write a storage group ACS routine that assigns a storage group that contains the volumes explicitly specified by the user
- Ensure all volumes explicitly specified by the user belong to the same storage group
- Ensure the requested space is available.

Assigning Storage Classes

A data set that is assigned a storage class is SMS managed, but a data set that is not assigned a storage class is non-SMS managed. You can assign storage classes either through the storage class ACS routine or by explicit specification. If the storage class ACS routine determines a storage class, it takes precedence over one that is explicitly specified by any of the following:

- JCL DD statements
- TSO/E ALLOCATE command
- DFDSS COPY and RESTORE commands
- Access method services ALLOCATE, DEFINE, and IMPORT commands
- Dynamic allocation requests
- ISPF/PDF data set allocation panels

The format for specifying a storage class on a JCL statement is:

STORCLAS=storage-class-name

The format for specifying a storage class on a TSO/E command is:

STORCLAS(storage-class-name)

The format for specifying a storage class on an access method services command is:

STORAGECLASS(-)

For information on determining storage classes with ACS routines. see Chapter 8, "Defining ACS Routines" on page 63.

Defining Additional Storage Classes

You can copy existing storage classes and modify them to create new storage classes by using the COPY line operator, which is explained in "Copying SMS Classes and Storage Groups" on page 100.

Chapter 7. Defining Data Classes

When allocating data sets, end users previously needed to specify many data set attributes. If they had many data sets with similar allocation requirements, they performed a lot of copying or repetitive typing. SMS simplifies the allocation of data sets by introducing data classes. This chapter describes data classes and shows you how to define them using the Data Class Application of ISMF.

Understanding Data Classes

A data class is a list of data set allocation attributes and their values. When end users create a data set and refer to a data class either explicitly (for example, through JCL) or implicitly (through ACS routines), SMS allocates the data set using the attribute values of its associated data class. For data class attributes, explicit specifications by end users override any parameters derived from the data class ACS routine. If SMS is active, a data class can be assigned to any DASD data set. For data sets that are not SMS managed. the system uses the allocation attribute values of the data class, but it does not save the data class name in the data set's catalog entry.

Not all attributes apply to every data set organization. When SMS allocates a data set, it uses only those data class attributes that have meaning for the given data set organization. SMS saves the data class name for each data set. The actual data class definitions reside in the SCDS. If you alter a data class definition, SMS applies the changes to any new data sets that use the data class after you activate the changed configuration. However, SMS does not retroactively apply your changes to previously created data sets.

Describing Data Classes

A data class definition contains identification and allocation information. To identify and refer to your data classes, you assign each one a unique name that contains from one to eight alphanumeric characters. Each data class maintains an owner ID that identifies the storage administrator who originally created or last modified the data class. Also, each data class contains an optional 120 character description field for describing its contents.

The data class allocation attributes match the keywords that you use to create data sets. The attributes contain space-related, access-related, and organizational information that you typically find on JCL DD statements, TSO/E ALLOCATE commands, access method services DEFINE commands, dynamic allocation requests, and ISPF/PDF panels.

Planning Data Classes

You should create a data class based on service level agreements. For example, you probably want to provide data classes for all data sets having a low-level qualifier of LIST, LISTING, OUTLIST, or LINKLIST probably belong to the same data class, because they are typical work data sets having similar allocation characteristics.

Before you actually define any data classes, gather information about the common types of data sets in your installation. You also need to determine if only the data class ACS routine can assign data classes to data sets, if end users can assign data classes to data sets, or if you want a combination of these two policies. If you intend to have only the data class ACS routine assign data classes, you need to develop methods to identify the data in your installation. However, if you allow only the data class ACS routine to assign data classes, you should be aware that users will need to override some data class attributes. For example, you will probably not want to code one data class for each possible amount of space that users will need. Finally, you need to identify the space requirements for some commonly used data sets.

By gathering useful installation information, you can relieve your end users from specifying all of the allocation attributes on allocation requests. They can use the data class that most closely matches their needs and explicitly change the few attributes that are unique to their data sets.

Defining Data Class Attributes

You can use ISMF to define your data classes by selecting option 4, DATA CLASS, from the ISMF Primary Option Menu for Storage Administrators. Figure 22 illustrates the Data Class Application Selection panel.

```
                         DATA CLASS APPLICATION SELECTION
COMMAND ===>

TO PERFORM DATA CLASS OPERATIONS, SPECIFY:

   CDS NAME         ===> 'SMS.SCDS1.SCDS'
                                   (1 to 44 character data set name or 'ACTIVE')

   DATA CLASS NAME ===> DATAV            (For Data Class List, fully or
                                          partially specified or * for all)

SELECT ONE OF THE FOLLOWING OPTIONS  ===> 3

   1   LIST    - Generate a list of Data Classes
   2   DISPLAY - Display a Data Class
   3   DEFINE  - Define a Data Class
   4   ALTER   - Alter a Data Class

IF OPTION 1 CHOSEN ABOVE,
   RESPECIFY SORT CRITERIA           ===> N   (Y or N)


USE ENTER TO PERFORM SELECTION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 22. Defining a Data Class

ISMF primes CDS NAME and DATA CLASS NAME with the most recently referenced values.

To define a data class, you must specify the name of an SCDS in CDS NAME, and you must provide a name in the DATA CLASS NAME field. (The default CDS NAME is the quoted word 'ACTIVE', which represents the currently active configuration, but you cannot define data classes to the 'ACTIVE' configuration.) Then select option 3, DEFINE, and press ENTER.

Defining Data Class Record and Space Attributes

Figure 23 shows the first of two pages for the Data Class Define panel. You can leave either page of the Data Class Define panel at any time without saving the data class by issuing the CANCEL command.

```
/                                                                    \
|                         DATA CLASS DEFINE               Page 1 of 2 |
|   COMMAND ===>                                                      |
|                                                                    |
|   SCDS NAME:       SMS.SCDS1.SCDS                                   |
|   DATA CLASS NAME: DATAV                                            |
|                                                                    |
|   TO DEFINE DATA CLASS, SPECIFY:                                    |
|                                                                    |
|     DESCRIPTION ===> Variable length record format for SCRIPT data sets, |
|                 ===> CLISTS, and program source code.              |
|     RECORG          ===>              (KS, ES, RR, LS or blank)     |
|     RECFM           ===> VB           (any valid RECFM combination or blank) |
|     LRECL           ===> 255          (1 to 32761 or blank)         |
|     KEYLEN          ===>              (0 to 255 or blank)           |
|     KEYOFF          ===>              (0 to 32760 or blank)         |
|     SPACE AVGREC    ===> U            (U, K, M or blank)            |
|           AVG VALUE ===> 255          (0 to 65535 or blank)         |
|           PRIMARY   ===> 5000         (0 to 999999 or blank)        |
|           SECONDARY ===> 5000         (0 to 999999 or blank)        |
|           DIRECTORY ===>              (0 to 999999 or blank)        |
|     RETPD OR EXPDT  ===>              (0 to 9999, yyyy/mm/dd or blank) |
|                                                                    |
|   USE ENTER TO PERFORM VERIFICATION; USE DOWN COMMAND TO VIEW NEXT PANEL; |
|   USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT. |
\                                                                    /
```

Figure 23. Defining Data Class Record and Space Attributes

The first page of the Data Class Define panel contains data class record and space attributes. SCDS NAME and DATA CLASS NAME are output fields that contain the SCDS and data class names you specified in the Data Class Application Selection panel. DESCRIPTION is an optional field of 120 characters where you can describe the data class. All of the remaining fields are optional and have a default value of blanks.

A blank indicates that no value is assigned to a parameter, and end users need to specify the value explicitly if is required and they intend to use the data class. Otherwise, their jobs might fail. When users override some data class attributes, they should ensure that the remaining attributes in the data class do not conflict with the explicitly specified attributes. For example:

```
LRECL < KEY OFFSET
DSORG = PO and 0 directory blocks.
```

See the *MVS/ESA Data Administration Guide* for more information. You can specify the following attributes on the first page of the Define Data Class panel:

RECORG   Specifies the data set organization, and it resembles the RECORG DD attribute. If you specify RECORG, you cannot specify the RECFM attribute. They are mutually exclusive. A blank value specifies either a physical sequential or a partitioned organization (non-VSAM data set).

RECFM   Specifies the data set record format. You can specify one of the following, with the option of appending either an A (ANSI carriage control) or an M (Machine carriage control):

- U—Undefined
- V—Variable
- VS—Variable spanned
- VB—Variable blocked
- VBS—Variable blocked spanned
- F—Fixed
- FS—Fixed standard
- FB—Fixed blocked
- FBS—Fixed blocked standard

If you specify RECFM, you cannot specify the RECORG attribute. They are mutually exclusive.

LRECL   Specifies the logical record length in bytes. If you leave RECORG blank, you can specify an LRECL from 1 to 32760 (or leave it blank). If you choose a RECORG of KS. ES, or RR. you can specify an LRECL from 1 to 32761 (or leave it blank). If RECORG is KS, LRECL must be greater than or equal to the value you specify for KEYLEN.

If RECORG is LS, LRECL is ignored and CISIZE is 4096. See the description of CISIZE DATA on page 58 for more information.

KEYLEN   Specifies the key length in bytes. To use this attribute, you must specify either KS or blank for RECORG. If RECORG is KS. KEYLEN represents the length of the KSDS key field and ranges from 1 to 255. or you can leave it blank. If RECORG is blank, KEYLEN ranges from 0 to 255, or you can leave it blank. For either value of RECORG, you must assign a value to KEYLEN that is less than or equal to LRECL.

KEYOFF   Specifies the displacement from the beginning of a record to the KSDS key field. It is only valid when RECORG is KS. KEYOFF can range from 0 to the value of (LRECL - KEYLEN).

SPACE   Specifies a request for space in bytes or multiples of bytes, eliminating track and cylinder space requests.

    AVGREC   Specifies a scaling factor for primary and secondary record allocations. It can have the following factors:

        U Multiplies the allocation quantity by 1.
        K Multiplies the allocation quantity by 1024.
        M Multiplies the allocation quantity by 1024^2 (= 1,048.576).

    AVG VALUE   Specifies the average length of each record.

    PRIMARY   Specifies the primary allocation quantity.

    SECONDARY   Specifies the secondary allocation quantity.

    DIRECTORY   Specifies the number of directory blocks for a PDS. It is only valid when the following are blank:

- RECORG
- KEYOFF
- IMBED

- REPLICATE
- CISIZE DATA
- % FREESPACE CI
- % FREESPACE CA
- SHAREOPTIONS XREGION
- SHAREOPTIONS XSYSTEM

Figure 24 illustrates the use of the space attributes.

```
If you specify the following values:

    AVGREC = K
    AVG VALUE = 80
    PRIMARY = 500
    SECONDARY = 100
    DIRECTORY = 100

you request 40000 KB of primary space, 8000 KB of secondary space, and
100 directory blocks.  JCL to accomplish the same task resembles:

//DD1  ...   SPACE=(80,(500,100,100)),AVGREC=K   ...
```

Figure 24. Space Attributes

RETPD
EXPDT RETPD and EXPDT specify how long a data set remains acces-
sible. RETPD specifies the retention period in days. and ranges
from 0 to 9999. EXPDT specifies the expiration date.

Defining Data Class Volume and VSAM Attributes

After specifying your attribute values. issue the DOWN command to view the
second page of the Data Class Define panel, which is shown in Figure 25 on
page 58.

```
┌─────────────────────────────────────────────────────────────────────────────┐
│                         DATA CLASS DEFINE                      Page 2 of 2    │
│  COMMAND ===>                                                                 │
│                                                                               │
│  SCDS NAME:        SMS.SCDS1.SCDS                                             │
│  DATA CLASS NAME: DATAV                                                       │
│                                                                               │
│  TO DEFINE DATA CLASS, SPECIFY:                                              │
│                                                                               │
│     VOLUME COUNT          ===> 1          (1 to 59 or blank)                  │
│     IMBED                 ===>            (Y, N or blank)                      │
│     REPLICATE             ===>            (Y, N or blank)                      │
│     CISIZE DATA           ===>            (1 to 32768 or blank)               │
│     % FREESPACE CI        ===>            (0 to 100 or blank)                 │
│               CA          ===>            (0 to 100 or blank)                 │
│     SHAREOPTIONS XREGION  ===>            (1 to 4 or blank)                   │
│                  XSYSTEM  ===>            (3, 4 or blank)                      │
│                                                                               │
│                                                                               │
│                                                                               │
│                                                                               │
│  USE ENTER TO PERFORM VERIFICATION; USE UP COMMAND TO VIEW PREVIOUS PANEL;    │
│  USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT. │
└─────────────────────────────────────────────────────────────────────────────┘
```

Figure 25. Defining Data Class Volume and VSAM Attributes

The second page of the Data Class Define panel contains volume and VSAM attributes for the data class. SCDS NAME and DATA CLASS NAME are output fields that contain the SCDS and data class names you specified on the Data Class Application Selection panel. All of the remaining fields are optional.

VOLUME COUNT Specifies the maximum number of DASD volumes that a data set can span. The default is 1.

IMBED Specifies whether the sequence set is to be written with the data component. If you specify Y, yes, the sequence set record for each control area is written on the first track adjacent to the control area, as many times as it fits. If you specify N, no, the sequence set record for each control area is written with the other index records. IMBED is only valid if RECORG is KS or blank. The default is blank.

REPLICATE Specifies whether each index record is to be written on a track as many times as it fits (Y) or only once (N). REPLICATE is only valid if RECORG is KS. The default is blank.

CISIZE DATA Specifies the control interval size for the data component of data sets having RECORG of KS, ES, or RR. It is calculated using the value specified in LRECL on page 1 of the Data Class Define panel. If you specify a RECORG of LS, CISIZE is automatically 4096, regardless of what you specify in the CISIZE DATA field. The default is blank.

% FREESPACE Specifies the percentage of space to be left free when a cluster is loaded. It also specifies the percentage of space to be left free after any split of control intervals (CI) or control areas (CA). It applies only to the data component. The FREESPACE attribute is only valid if RECORG is KS or blank. The default is blank.

SHAREOPTIONS Specifies how end users can share a component or cluster.

XREGION specifies the amount of sharing allowed among regions within the same system, or within multiple systems using GRS. You can specify the following values:

1 Any number of users can read the data set at one time, *or* only one can write to it.

2 Any number of users can read the data set at one time, *and* only one can write to it.

3 Any number of users can share the data set, and each is responsible for maintaining read and write integrity.

4 Any number of users can share the data set, and buffers used for direct processing are refreshed for each request.

XSYSTEM specifies the amount of sharing allowed among systems. You can specify the following values:

3 Any number of users can share the data set, and each is responsible for maintaining read and write integrity.

4 Any number of users can share the data set, and buffers used for direct processing are refreshed for each request.

If values are not specified explicitly or in the data class, VSAM defaults are used.

Figure 26 summarizes which data class attributes apply to each data organization (RECORG).

Figure 26. Applying Data Class Attributes to RECORG

Attribute	Record Organization (RECORG)				
	Blank	KS	ES	RR	LS
LRECL	x	x	x	x	
RECFM	x				
KEYLEN	x	x			
KEYOFF	x	x			
VOLUME COUNT	x	x	x	x	x
SPACE	x	x	x	x	x
CISIZE	x	x	x	x	
IMBED	x	x			
REPLICATE	x	x			
FREESPACE	x	x			
SHAREOPTIONS	x	x	x	x	x
RETPD or					
EXPDT	x	x	x	x	x

See the *MVS/ESA Integrated Catalog Adminstration: Access Method Services Reference*, the *MVS/ESA VSAM Administration Guide*, and the *MVS/ESA JCL Reference* for more information.

Assigning Data Classes

You can define an ACS routine to determine data classes, or end users can explicitly specify a data class name on the following:

- JCL DD statements
- TSO/E ALLOCATE commands
- Access method services ALLOCATE and DEFINE commands
- Dynamic allocation requests
- ISPF/PDF data set allocation panels.

Note that end users can explicitly specify data set attributes, and these specifications take precedence over the data set attributes assigned through the data class ACS routine.

The format for specifying a data class on a JCL statement is:

```
DATACLAS=data-class-name
```

The format for specifying a data class on a TSO/E command is:

```
DATACLAS(data-class-name)
```

The format for specifying a data class on an access method service command is:

```
DATACLASS(-)
```

Processing of Data Class Attributes in JCL

The order of precedence for data class attributes in JCL is as follows:

1. Explicit specifications
2. LIKE and REFDD keywords
3. Data class definitions (explicit or derived) in REFDD statement
4. Data class definitions in the referencing DD.

A data class does not need to be a self-contained, complete data set organization. You can partially define the data set attributes in the data class definition (as a base) and then explicitly specify the remaining attributes. However, the merging of all attributes according to the order specified above must result in a valid data organization. In the example below,

```
//DD1   DD    ...,DATACLAS=DC2,RECORG=ES,...
//DD2   DD    ...,DATACLAS=DC1,LRECL=180,REFDD=DD1,...
```

The LRECL of 180 in DD2 is used first, and then the RECORG of ES in DD1 is used, regardless of the values specified in either the DC1 or the DC2 data class. For the remaining attributes that are not explicitly specified on the DD statements, SMS uses the values defined in the data class definition of DC2 and then DC1.

In this next example, the attributes of the data set referenced by the LIKE keyword are used after all other explicit specification but before data class DC3 attributes:

```
//DD3   DD    ...,DATACLAS=DC3,LRECL=180,LIKE=SAMPLE.DATA,....
```

Here, an LRECL of 180 is used. Then, SMS uses the DSCB information from the SAMPLE.DATA data set. Finally, the remaining attribute values are drawn from the DC3 data class.

A final example illustrates the use of REFDD:

```
//DD4  DD    ...,DATACLAS=DC4,...
//DD5  DD    ...,DSN=DS1,REFDD=DD4,DATACLAS=DC5,LRECL=180,...
```

In this example, the REFDD keyword specifies that the explicit attributes on the DD4 JCL statement are to be used second since explicit attributes on the DD4 JCL statement are used after the explicitly specified attributes. Next, the attributes from data class DC4 referenced in DD4 should be used. The remaining attributes are taken from data class DC5.

Note that the LIKE and REFDD keywords are mutually exclusive.

For information on determining data classes through ACS routines, see Chapter 8, "Defining ACS Routines" on page 63. For information on IDCAMS DEFINE defaults and data class defaults, see the *MVS/ESA Integrated Catalog Adminstration: Access Method Services Reference*.

Defining Additional Data Classes

You can copy existing data classes and modify them to create new ones by using the COPY line operator from the LIST panel, which is explained in "Copying SMS Classes and Storage Groups" on page 100.
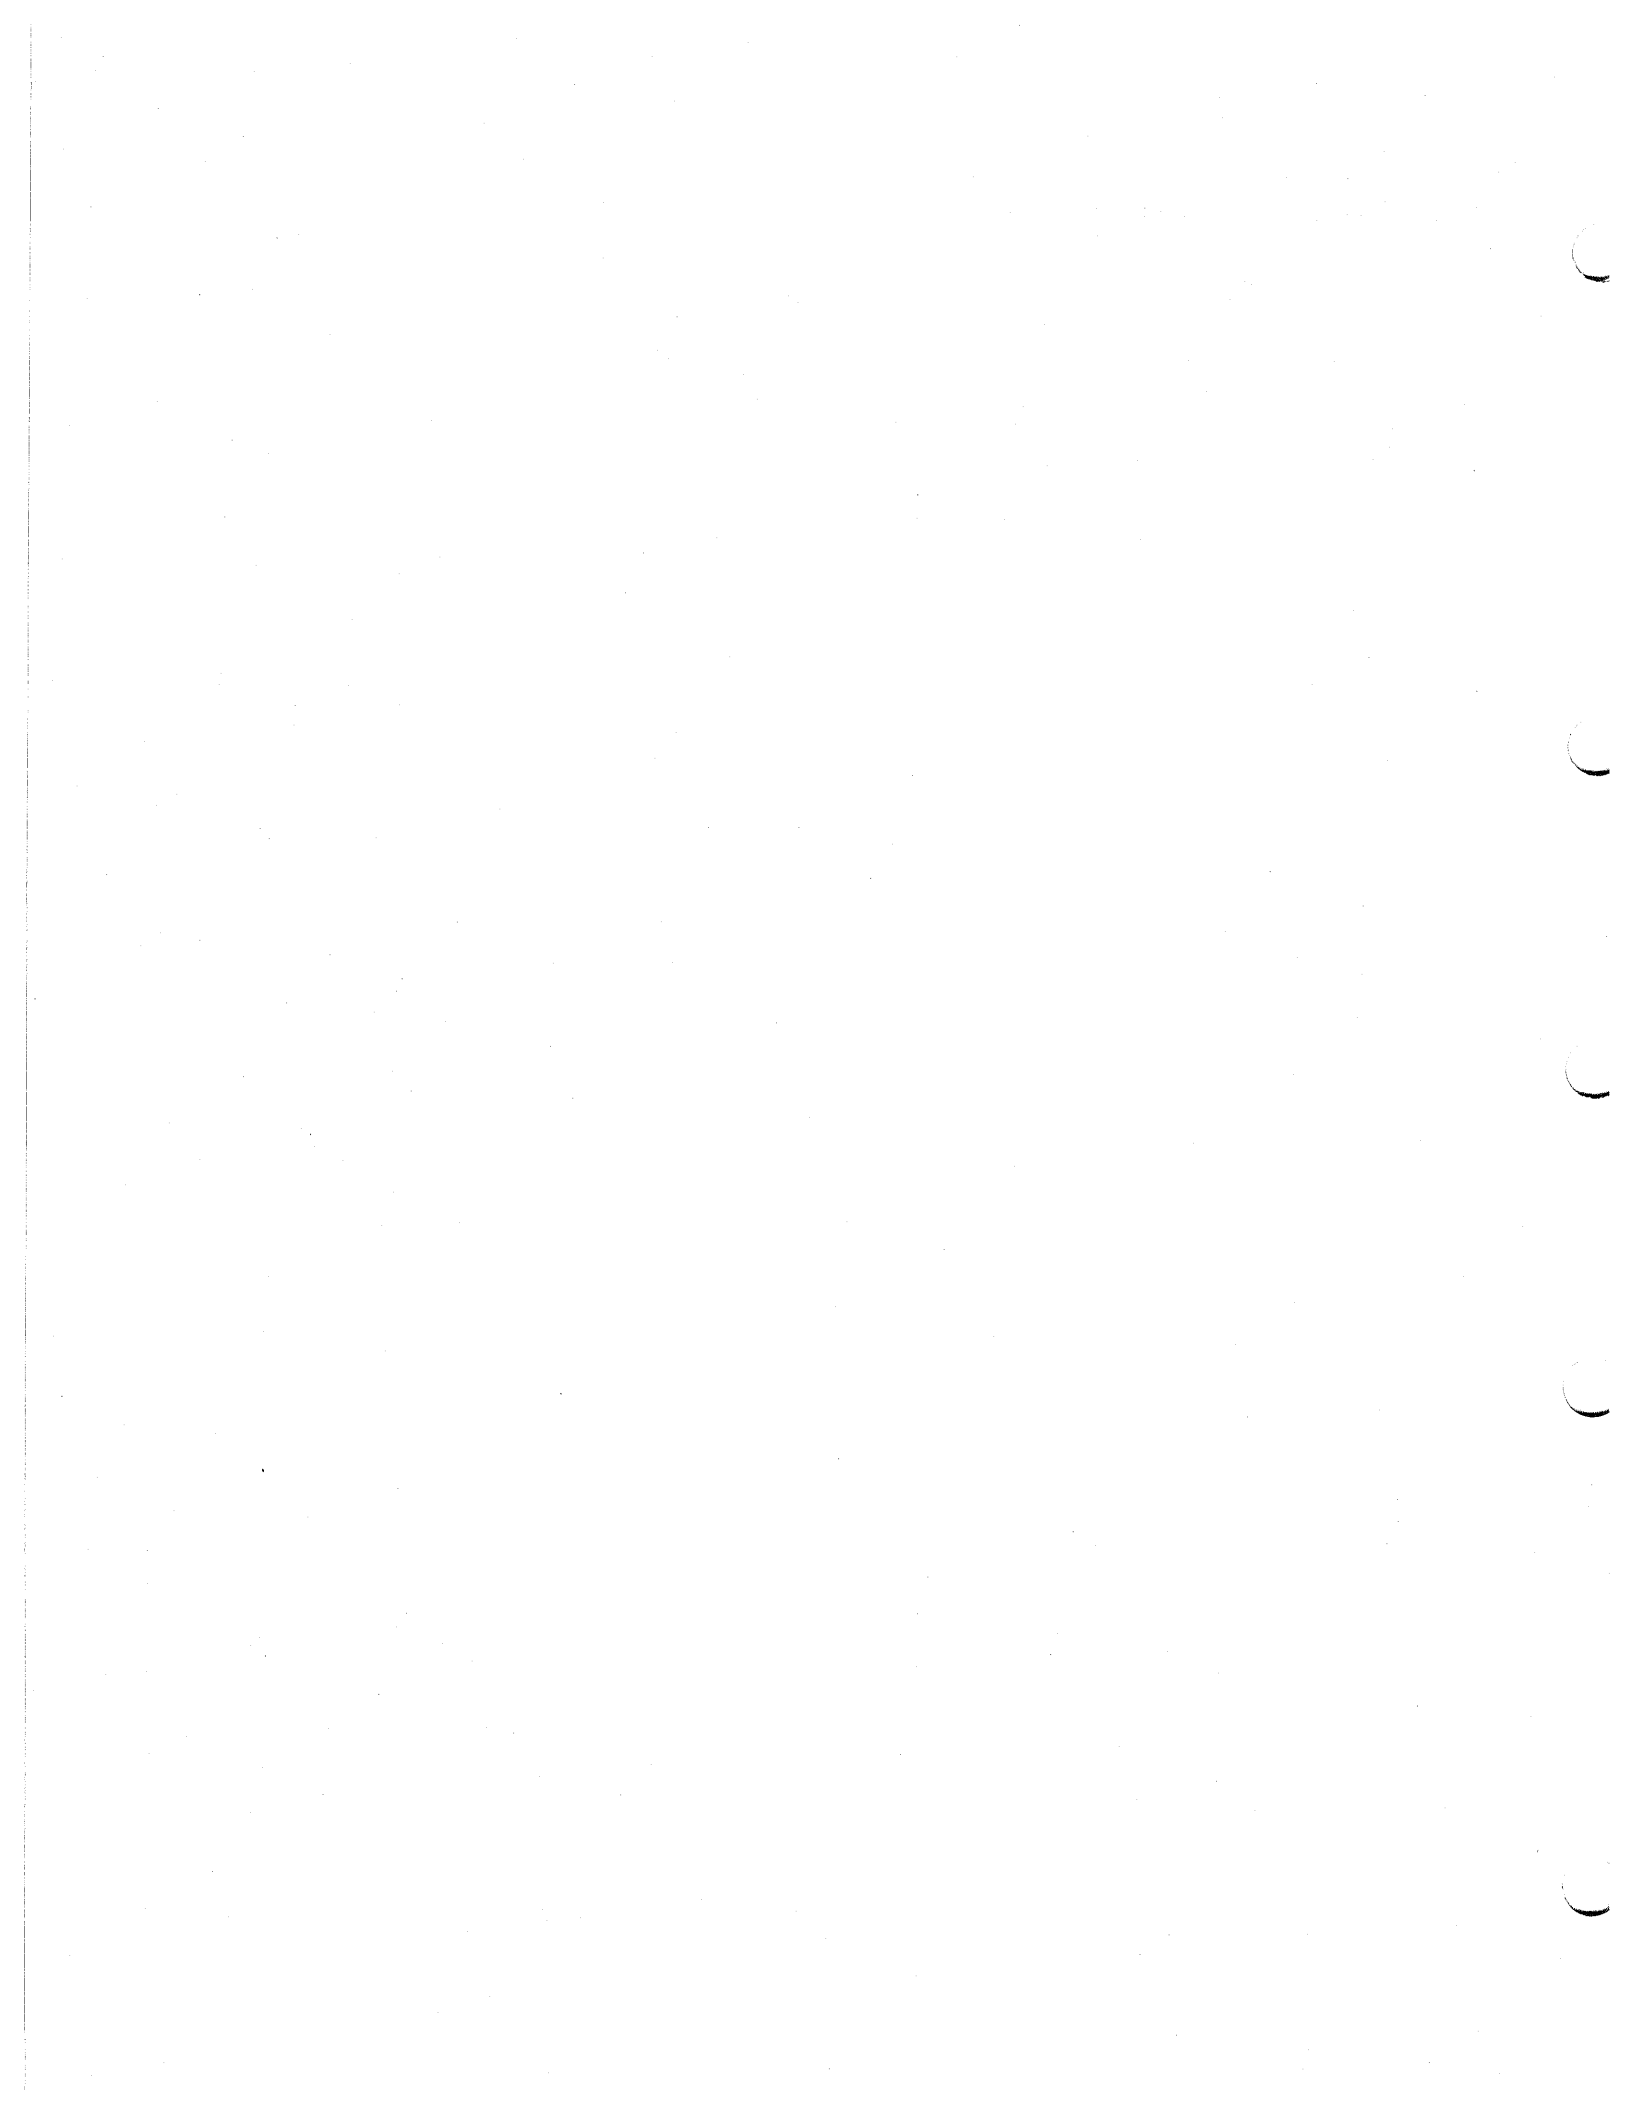
Chapter 8. Defining ACS Routines

ACS routines can be used to determine the SMS classes and storage groups for data sets in an SMS complex. For storage administrators, ACS routines automate and centralize the process of determining SMS classes and storage groups. They also facilitate the conversion of data sets to an SMS environment. This chapter explains how to write ACS routines for an SMS configuration using the Automatic Class Selection Application of ISMF.

Chapter 13, "ACS Language Reference" on page 117 lists the rules for programming in the ACS language.

Understanding ACS Routines

Through ISMF, you can create and maintain as many as four ACS routines in an SCDS, one for each type of SMS class and storage group. After you have activated an SMS configuration, SMS executes the ACS routines for the following operations:

- JCL DD statements (NEW, MOD)
- Dynamic allocation requests (NEW, MOD)
- DFDSS COPY, RESTORE, and CONVERTV commands
- DFHSM RECALL and RECOVER
- Access method services ALLOCATE, DEFINE, and IMPORT commands.

As a storage administrator, you write ACS routines using the ACS programming language, a high-level programming language. The language follows a logical, procedural flow of execution that consists mainly of filtering criteria, IF/THEN statements, and SELECT/WHEN statements. Using these relational statements, ACS routines determine SMS classes and storage groups according to allocation parameters, data set sizes, and other variables.

Note: Only the storage group ACS routine is required. The other routines are optional.

Creating ACS Routines

You must allocate either a sequential data set or a member of a partitioned data set (PDS) for each of the ACS routines that you intend to write. As a general guideline, use an LRECL of 80. After allocating the data sets, select option 7, Automatic Class Selection Application, from the ISMF Primary Option Menu and press ENTER. You get the Automatic Class Selection Application Selection panel shown in Figure 27 on page 64.

```
                          ACS APPLICATION SELECTION
COMMAND ===>

SELECT ONE OF THE FOLLOWING OPTIONS ===> 1

    1  EDIT           - Edit ACS Routine source code
    2  TRANSLATE      - Translate ACS Routines to ACS Object Form
    3  VALIDATE       - Validate ACS Routines Against Storage Constructs
    4  TEST           - Define/Alter Test Cases and Test ACS Routines
    5  DISPLAY        - Display ACS Object Information
    6  DELETE         - Delete an ACS Object from a Source Control Data Set




IF OPTION 5 CHOSEN ABOVE, SPECIFY:

  CDS NAME  ===> 'SMS.SCDS1.SCDS'
                            (1 to 44 character data set name or 'ACTIVE')



USE ENTER TO PERFORM SELECTION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 27. Writing an ACS Routine

Select option 1, EDIT, and press ENTER to write an ACS routine. ISPF/PDF displays the Edit Entry panel shown in Figure 28.

```
--------------------------- EDIT - ENTRY PANEL ---------------------------
COMMAND ===>

ISPF LIBRARY:
    PROJECT ===>
    GROUP   ===>          ===>          ===>          ===>
    TYPE    ===>
    MEMBER  ===>              (Blank or pattern for member selection list)

OTHER PARTITIONED OR SEQUENTIAL DATA SET:
    DATA SET NAME  ===>
    VOLUME SERIAL  ===>      (If not cataloged)

DATA SET PASSWORD ===>       (If password protected)

PROFILE NAME      ===>       (Blank defaults to data set type)

INITIAL MACRO     ===>       LOCK      ===> YES    (YES, NO or NEVER)

FORMAT NAME       ===>       MIXED MODE ===> NO    (YES or NO)
```

Figure 28. Invoking the ISPF/PDF Editor

After specifying your values, press ENTER to summon the ISPF/PDF Editor. Then enter the source code for the ACS routine. After entering the routine, issue the END command to save it and return to the ISPF/PDF Edit Entry panel. Enter END again to return to the Automatic Class Selection Application Selection panel.

Translating ACS Routines

After creating an ACS routine, you must translate it into executable form. The translation process checks your source code for syntactic and semantic errors, generates an object table if no errors exist, and places the object table into an SCDS. If the ACS routine that you are translating already exists in the SCDS, the new object table replaces the existing object table.

To translate an ACS routine, select option 2, TRANSLATE, from the Automatic Class Selection Application Selection panel. ISMF displays the Translate ACS Routines panel shown in Figure 29.

```
                          TRANSLATE ACS ROUTINES
COMMAND ===>

TO PERFORM ACS TRANSLATION, SPECIFY:

  SCDS NAME            ===> 'SMS.SCDS1.SCDS'
                                (1 to 44 character data set name)

  ACS SOURCE DATA SET ===> 'SCDS1.STORGRP'
                                (1 to 44 character data set name)

  ACS SOURCE MEMBER   ===>       (1 to 8 characters)


  LISTING DATA SET    ===> 'STORGRP.LISTING'
                                (1 to 44 character data set name)




USE ENTER TO PERFORM ACS TRANSLATION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 29. Translating an ACS Routine

A successful translation places the ACS routine object table into the SCDS that you specify in SCDS NAME. It is a required field that is primed with the last SCDS that you have referenced or translated into. The default is blanks.

Specify the name of the data set containing the ACS routine that you want to translate in ACS SOURCE DATA SET. It is a required field that is primed with the last used value. The default is blanks.

ACS SOURCE MEMBER is the name of the member of the source library that contains the ACS routine. It is required only if the ACS SOURCE DATA SET is a partitioned data set. The field is primed with the last used value, and the default is blanks.

Specify the name of a sequential data set to contain the translation results in LISTING DATA SET. If you specify a data set that already exists, the translation process replaces the existing data set contents with the results of the translation. If you specify a new data set, the translation process allocates space for it. If you leave this field blank, which is the default, you receive the results of

the validation but you do not get a listing. After specifying your values, press
ENTER to perform the translation.

BROWSING the Results of a Translation

If you specified a data set name in the LISTING DATA SET field, you get the
Browse panel shown in Figure 30 after validation.

```
 BROWSE -- STORGRP.LISTING------------------------- LINE 00000000 COL 001 080
 COMMAND ===>                                            SCROLL ===> PAGE
 ****************************** TOP OF DATA ********************************
 ACS TRANSLATOR ***** TIME 09:45:32 DATE 11/19/1987 PAGE 0001 *****


 SCDS NAME:              SMS.SCDS1.SCDS

 ACS SOURCE DATA SET:    'SCDS1.STORGRP'
 ACS SOURCE MEMBER:



 0001      PROC STORGRP
 0002        SELECT
 0003          WHEN(&STORCLAS = DB*)
 0004            SET &STORGRP = 'DATABASE'
 0005          WHEN(&MAXSIZE > 9999999KB)
 0006            SET &STORGRP = 'LARGE'
 0007          WHEN(&DSTYPE = 'TEMP')
 0008            SET &STORGRP = 'VIO','PRIMARY'
 0009          OTHERWISE
 0010            SET &STORGRP = 'PRIMARY'
 0011        END
 0012      END


 TRANSLATION RETURN CODE:  0000
 ****************************** BOTTOM OF DATA ****************************
```

Figure 30. BROWSING the Results of an ACS Translation

The Browse panel displays the SCDS name, the ACS routine source data set
name, an output listing of the ACS routine, and a translation return code. SCDS
NAME, ACS SOURCE DATA SET, and ACS SOURCE MEMBER reflect the values
that you specified on the Translate ACS Routines panel. The output listing con-
tains the source code of the ACS routine and diagnostic messages. See
MVS/ESA Message Library: System Messages for an explanation of the diag-
nostic messages.

The TRANSLATION RETURN CODE displays one of the following codes:

0000 Successful translation.

0012 Unsuccessful translation. The ACS routine contains one or more
semantic or syntactic errors. No object table was created and no
updates were made to the SCDS.

0020 Internal error in translator.

When you exit from this panel, you get the Output Listing Disposition panel
shown in Figure 31 on page 67.

```
                         OUTPUT LISTING DISPOSITION      TRANSLATION SUCCESSFUL
COMMAND ===>

LISTING DATA SET:  STORGRP.LISTING

SPECIFY OUTPUT LISTING DISPOSITION:


   PRINT OUTPUT LISTING          ===> N       (Y or N)



   DELETE OUTPUT LISTING         ===> N       (Y or N)






USE ENTER TO PERFORM SELECTION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 31. Generating an ACS Translation Listing

The short message area displays the results of the translation. Specify Y, yes, in PRINT OUTPUT LISTING to submit a batch job to print the data set. Specify Y, yes, in DELETE OUTPUT LISTING to delete the data set after it is printed. Both fields are required, and they are primed with N, no. Figure 32 summarizes the possible outcomes for each combination of Output Listing Disposition values.

Figure 32. Results of Output Listing Disposition			
Print	**Delete**	**Resulting Action**	**Type of Job**
Yes	Yes	Print dataset DISP=(OLD,DELETE)	Batch
Yes	No	Print dataset DISP=(OLD,KEEP)	Batch
No	Yes	Delete dataset	Foreground
No	No	No action	N/A

Issue the END command to return to the Translate ACS Routines panel.

BROWSING the Results of an Unsuccessful Translation

If you specified a data set name in error on the LISTING DATA SET field, you get the Browse panel shown in Figure 33 on page 68.

```
  BROWSE -- STORGRP.LISTING-------------------------- LINE 00000000 COL 001 080
  COMMAND ===>                                              SCROLL ===> PAGE
  ******************************* TOP OF DATA  ***********************************
  ACS TRANSLATOR ***** TIME 11:47:09 DATE 09/22/1988 PAGE 0001 ****

  SCDS NAME:              SMS.SCDS.SCDS1
  ACS SOURCE DATA SET:    'SCDS1.STORGRP'
  ACS SOURCE MEMBER:

  0001       PROC STORGRP
  0002         SELECT
  0003           WHEN(&STORCLAS = DB*)
  0004             SET &STORGRP = 'DATABASE'
  0005           WHEN(&MAXSIZE > 9999999KB)
  *******      IGD03172I MISSING END KEYWORD FOR A SELECT STATEMENT
  *******      IGD03116I UNEXPECTED STATEMENT TYPE WHEN ENCOUNTERED - SUBSEQUENT
                         STATEMENTS MAY BE IGNORED
  0006             SET &STORGRP = 'LARGE'
  0007           WHEN &DSTYPE = 'TEMP'
  0008             SET &STORGRP = 'VIO','PRIMARY'
  *******      IGD03222I INVALID ASSIGNMENT OR MISSING LITERAL IN SET STATEMENT
  0009           OTHERWISE
  0010             SET &STORGRP = 'PRIMARY'
  0011         END
  0012       END

  TRANSLATION RETURN CODE:  0012
  *************************** BOTTOM OF DATA *****************************
```

Figure 33. BROWSING the Results of an Unsuccessful ACS Translation

Page down to see a display of the error messages.

Validating ACS Routines or an Entire SCDS

SMS validates the entire configuration when you modify an SCDS, but you can
validate the individual ACS routines of an SMS configuration after successfully
translating them. You should also validate the entire configuration yourself so
that you can see any error messages that result. Separate validation of the
ACS routines does not produce all of the possible messages.

Validating an ACS Routine

To validate an ACS routine, select option 3, VALIDATE on the Automatic Class
Selection Application Selection panel. ISMF displays the Validate ACS Routines
or Entire SCDS panel shown in Figure 34 on page 69.

```
┌─────────────────────────────────────────────────────────────────────────┐
│                    VALIDATE ACS ROUTINES OR ENTIRE SCDS                   │
│  COMMAND ===>                                                             │
│                                                                           │
│                                                                           │
│  TO PERFORM VALIDATION, SPECIFY:                                          │
│                                                                           │
│    SCDS NAME          ===> 'SMS.SCDS1.SCDS'                               │
│                                  (1 to 44 character data set name)        │
│                                                                           │
│    ACS ROUTINE TYPE ===> SG          (DC=Data Class, MC=Management Class, │
│                                       SC=Storage Class, SG=Storage Group, │
│                                       *=entire SCDS)                      │
│                                                                           │
│    LISTING DATA SET ===> 'SGVAL.LISTING'                                  │
│                                  (1 to 44 character data set name)        │
│                                                                           │
│                                                                           │
│                                                                           │
│                                                                           │
│                                                                           │
│  USE ENTER TO PERFORM VALIDATION;                                         │
│  USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.                      │
└─────────────────────────────────────────────────────────────────────────┘
```

Figure 34. Validating an ACS Routine

SCDS NAME identifies the SCDS containing the ACS routine you are validating.
It is a required field that is primed with the last used value. The default is
blanks.

ACS ROUTINE TYPE identifies which of the four ACS routine types you are vali-
dating. When validating an individual ACS routine, the valid values are: DC,
SC, MC, and SG. It is a required field that is primed with the last used value.
The default is an asterisk, *, which specifies that you want to validate the entire
SMS configuration in the SCDS, and not just one of the ACS routines.

LISTING DATA SET identifies the name of a sequential data set to contain the
results of the validation. If you specify a data set that already exists, the vali-
dation process replaces the existing data set contents with the results of the
validation. If you specify a new data set, the validation process allocates space
for it. If you leave this field blank, which is the default, you receive the result of
the validation, but you do not get a listing.

After specifying your values, press ENTER to perform the validation. Validation
fails if the following conditions are not satisfied (except where it is noted that
only a warning is issued):

- For the storage group ACS routine:

    All of the defined pool storage groups must be possible outcomes of the
    routine. Assume you have five storage groups, four of which are type
    'pool'. If only three different pool storage groups are possible outcomes
    of the storage group ACS routine, the validation process fails. The
    unreferenced fourth pool storage group represents a wasted resource,
    because SMS will never allocate a data set to that storage group.

    No dummy storage group can result from the storage group ACS
    routine.

All VIO storage groups are possible results of the storage group ACS routine.

All storage groups that are possible outcomes of the storage group ACS routine must exist.

- For the management class ACS routine:

  All management classes that are possible outcomes of the management class ACS routine must exist.

- For the storage class ACS routine:

  All storage classes that are possible outcomes of the storage class ACS routine must exist.

- For the data class ACS routine:

  All data classes that are possible outcomes of the data class ACS routine must exist.

If an ACS routine uses the SET statement to assign an SMS class or storage group that does not exist, validation fails. If an ACS routine references an SMS storage class, management class, or data class that does not exist, SMS issues a warning. For example, you might delete a storage class sometime in the future. However, if you reference the deleted storage class in the storage class ACS routine of an SCDS, SMS issues a warning when you attempt to validate the SCDS. If no other errors exist, the SCDS becomes valid despite the warning. This allows you to check for the deleted storage class and replace it with an existing storage class within the storage class ACS routine.

BROWSING the Results of a Validation

If you specified a data set name in the LISTING DATA SET field, you get the Browse panel shown in Figure 35 after translation.

```
 BROWSE -- SGVAL.LISTING --------------------------- LINE 00000000 COL 001 080
 COMMAND ===>                                             SCROLL ===> PAGE
 ****************************** TOP OF DATA *********************************
                         VALIDATION RESULTS

 VALIDATION RESULT:  ERRORS DETECTED

 SCDS NAME:          SMS.SCDS1.SCDS

 ACS ROUTINE TYPE:   SG
 DATE OF VALIDATION: 1987/11/19
 TIME OF VALIDATION: 10:01

 IGD06025I THE STORAGE GROUP ACS ROUTINE SETS NON-EXISTENT VALUE DATABASE
 IGD06025I THE STORAGE GROUP ACS ROUTINE SETS NON-EXISTENT VALUE LARGE
 IGD06025I THE STORAGE GROUP ACS ROUTINE SETS NON-EXISTENT VALUE VIO
 IGD06025I THE STORAGE GROUP ACS ROUTINE SETS NON-EXISTENT VALUE PRIMARY
 IGD06025I THE STORAGE GROUP ACS ROUTINE SETS NON-EXISTENT VALUE PRIMARY
 ****************************** BOTTOM OF DATA *****************************
```

Figure 35. BROWSING the Results of an ACS Routine Validation

VALIDATION RESULT is either **VALIDATION SUCCESSFUL** or **ERRORS DETECTED**. SCDS NAME and ACS ROUTINE TYPE reflect the values you specified on the Validate ACS Routines or Entire SCDS panel. Browse also provides

the date and time of validation. See *MVS/ESA Message Library: System Messages* for an explanation of the diagnostic messages.

When you leave Browse, you receive the Output Listing Disposition panel. See page 67 for information about printing and deleting.

Validating an Entire SCDS

You can validate an entire SCDS using the ISMF VALIDATE command. To use the command, type VALIDATE on the command line of the CDS Application Selection panel (shown in Figure 6 on page 19) and press ENTER. You receive the Validate ACS Routines or Entire SCDS panel shown in Figure 36.

```
                        VALIDATE ACS ROUTINES OR ENTIRE SCDS
 COMMAND ===>


 TO PERFORM VALIDATION, SPECIFY:
   SCDS NAME          ===> 'SMS.SCDS1.SCDS'
                                        (1 to 44 character data set name)
   ACS ROUTINE TYPE ===> *              (DC=Data Class, MC=Management Class,
                                         SC=Storage Class, SG=Storage Group,
                                         *=entire SCDS)

   LISTING DATA SET ===>
                                        (1 to 44 character data set name)




 USE ENTER TO PERFORM VALIDATION
 USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 36. Validating an SCDS

The SCDS NAME field identifies the data set you want to validate. It is a required field that is primed with the last SCDS you have referenced or validated. You should place an asterisk (*) in the ACS ROUTINE TYPE field, because you will be validating an entire SCDS. You can save a listing of the validation results by specifying a sequential data set name in the LISTING DATA SET field. If you specify a data set that already exists, you will replace the old copy with your new listing. If you specify a new data set, ISMF will allocate space for it. If you leave this field blank, you will not generate a listing.

Validation fails if any of the following conditions are not satisfied (except where it is noted that only a warning is issued).

- Base configuration information must exist. If you specify a default management class, you need to have defined the management class.

- Each possible outcome of the data class, storage class, management class, and storage group ACS routines must have a corresponding definition in the SCDS. For example, if you have five unique storage classes that can be determined by the storage class ACS routine, then you must have five corresponding storage class definitions in the SCDS.

- All classes and storage groups referenced in all ACS routines exist (issues a warning only).

- The storage group ACS routine exists.

- At least one pool type storage group exists.

- All pool and dummy storage groups have at least one volume.

- At least one storage class exists.

- All pool and VIO storage groups are possible outcomes.

- No dummy storage group set by storage group ACS.

Testing ACS Routines

After completing your ACS routines, you can write and execute test cases using the Automatic Class Selection Application of ISMF. After testing the individual routines of a new or modified configuration, you can activate it with greater confidence. ACS installation exits are not invoked during ACS routine testing.

Creating ACS Test Cases

You need to allocate a PDS for the ACS test cases. After allocating the data set, go to the ISMF Primary Option Menu and select option 7, Automatic Class Selection Application. To specify test cases, select option 4, TEST, from the Automatic Class Selection Application Selection panel. When you press ENTER, you get the ACS Test Selection panel shown in Figure 37.

```
                          ACS TEST SELECTION
COMMAND ===>

SELECT ONE OF THE FOLLOWING OPTIONS ===> 1

    1  DEFINE        - Define an ACS Test Case
    2  ALTER         - Alter an ACS Test Case
    3  TEST          - Test ACS Routines


IF OPTION 1 OR 2 CHOSEN ABOVE, SPECIFY:

  ACS TEST LIBRARY ===> SCDS1.ACSTEST.LIB
  ACS TEST MEMBER  ===> TEST1






USE ENTER TO PERFORM SELECTION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 37. Creating ACS Test Cases

The ACS TEST LIBRARY field identifies the data set that contains the ACS test case. The ACS TEST MEMBER field identifies which member of the ACS test library contains the desired test case. You can specify one test case per member. After specifying a test library and member, select option 1 and press ENTER to specify test criteria. You receive the first page of the ACS Test Case

Define panel, which is shown in Figure 38. The second page is shown in Figure 39, and the third page is shown in Figure 40 on page 74. This test case checks to see that the data base storage group is a possible outcome of the storage class ACS routine.

```
                          ACS TEST CASE DEFINE                  Page 1 of 3
COMMAND ===>

ACS TEST LIBRARY: SCDS1.ACSTEST.LIB
ACS TEST MEMBER:  TEST1

TO DEFINE ACS TEST CASE, SPECIFY:

  DESCRIPTION ===> Make sure data base storage group
              ===> is a possible outcome.


  DSN      ===>
  DD       ===>         RETPD    ===>         MGMTCLAS     ===>
  DSORG    ===>         XMODE    ===>         STORCLAS     ===> DBBASE
  RECORG   ===>         ACSENVIR ===>         DEF_DATACLAS ===>
  DSTYPE   ===>         SIZE     ===>         DEF_MGMTCLAS ===>
  DSOWNER  ===>         MAXSIZE  ===>         DEF_STORCLAS ===>
  EXPDT    ===>         DATACLAS ===>         DSNTYPE      ===>



  USE ENTER TO PERFORM VERIFICATION; USE DOWN COMMAND TO VIEW NEXT PANEL;
  USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 38. Creating ACS Test Cases - Page 1

```
                          ACS TEST CASE DEFINE                  Page 2 of 3
COMMAND ===>

ACS TEST LIBRARY: SCDS1.ACSTEST.LIB
ACS TEST MEMBER:  TEST1

TO DEFINE ACS TEST CASE, SPECIFY:

  JOB       ===>       PGM    ===>       VOL   ===>
  GROUP     ===>       APPLIC ===>       UNIT  ===>
  USER      ===>       NVOL   ===>       MSVGP ===>
  ACCT_JOB  ===>
            ===>
            ===>
  ACCT_STEP ===>
            ===>
            ===>







  USE ENTER TO PERFORM VERIFICATION; USE UP/DOWN COMMAND TO VIEW OTHER PANELS;
  USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 39. Creating ACS Test Cases - Page 2

```
/                                                                              \
              ACS TEST CASE VOLUME SERIAL SPECIFICATION          Page 3 of 3
   COMMAND ===>

   ACS TEST LIBRARY: SCDS1.ACSTEST.LIB
   ACS TEST MEMBER:   TEST1

   TO COMPLETE ACS TEST CASE DEFINE, SPECIFY:

     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>
     VOLSER ===>          ===>         ===>         ===>         ===>

     VOLSER ===>          ===>         ===>




   USE ENTER TO PERFORM VERIFICATION; USE UP COMMAND TO VIEW PREVIOUS PANEL;
   USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT.
\                                                                              /
```

Figure 40. Creating ACS Test Cases - Page 3

On these three panels, you can specify test values that correspond to the ACS
variables listed in Chapter 13, "ACS Language Reference" on page 117. Use
the END command to save your test values and return to the ACS Test
Selection panel. In the Test Selection panel, you can add more test case
members to the same test library, or you can build another test library.

Running ACS Test Cases

To run the test cases in a test library, specify the library name and select
option 3 in the ACS Test Selection panel (Figure 37 on page 72). You will
receive the Test ACS Routines panel shown in Figure 41 on page 75.

```
                             TEST ACS ROUTINES
  COMMAND ===>

  TO PERFORM ACS TESTING, SPECIFY:

    CDS NAME         ===> 'SMS.SCDS1.SCDS'
                                   (1 to 44 character data set name or 'ACTIVE')

    ACS TEST LIBRARY ===> SCDS1.ACSTEST.LIB
    ACS TEST MEMBER  ===> TEST1    (fully or partially specified
                                    or * for all members)
    LISTING DATA SET ===> 'TEST1.LISTING'
                                   (1 to 44 character data set name or BLANK)

  SELECT WHICH ACS ROUTINES TO TEST:

    DC ===> N (Y/N)    SC ===> N (Y/N)   MC ===> N (Y/N)    SG ===> Y (Y/N)




  USE ENTER TO PERFORM VERIFICATION AND TESTING;
  USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 41. Testing an ACS Routine

In this panel, you indicate which routines you want to test. Specifying an
asterisk, (*), in the ACS TEST MEMBER field is convenient for running all test
cases in the PDS library. You can create a listing of the results of the testing
by specifying a data set name in the LISTING DATA SET field. Leaving this field
blank prevents the creation of a listing.

You can now test the routines by pressing ENTER. If you specified a LISTING
DATA SET, the results of the testing will display in the PDF Browse panel shown
in Figure 42.

```
  BROWSE -- TEST1.LISTING -------------------------- LINE 00000000 COL 001 080
  COMMAND ===>                                              SCROLL ===> PAGE
  ******************************* TOP OF DATA  ********************************
                            ACS TESTING RESULTS

  CDS NAME          : SMS.SCDS1.SCDS

  ACS ROUTINE TYPES: SG

  ACS TEST LIBRARY : SCDS1.ACSTEST.LIB

    ACS TEST
    MEMBER       EXIT CODE   RESULTS
    ---------    ----------  ------------------------------------
    TEST1                 0  SG = DATABASE


  ACS TESTING RC:  00
  ***************************** BOTTOM OF DATA *******************************
```

Figure 42. Testing an ACS Routine

After examining the results, you can issue the END command. You will see the ACS Output Listing Disposition, where you can specify whether or not to keep the output listing.

ACS Routine Environments

Depending on the environment, SMS invokes some or all of the ACS routines in the following order:

1. Data class
2. Storage class
3. Management class
4. Storage group.

JCL DD Statement, (Batch), and Dynamic Allocation: For this environment (&ACSENVIR = 'ALLOC'), the data class ACS routine and then the storage class ACS routine are executed. If the storage class is not null, the management class ACS routine and then the storage group ACS routine are executed.

However, (VOL = REF) is an exception for the following cases:

- Data class ACS routine
- Storage class from referenced data set if the referenced data set has a storage class assigned to it
- Management class ACS routine
- Storage group from referenced data set.

Access Method Services: For the ALLOCATE and DEFINE commands (&ACSENVIR = 'ALLOC'), the data class ACS routine and then the storage class ACS routine are executed. If the storage class is not null, the management class ACS routine and then the storage group ACS routine are executed.

For the IMPORT command, the storage class ACS routine is executed first. If the storage class is not null, the management class ACS routine and then the storage group ACS routine are executed.

DFHSM: For the RECALL and RECOVER environments (&ACSENVIR = 'RECALL' and &ACSENVIR = 'RECOVER'), the storage class ACS routine is executed first. If the storage class is not null, the management class ACS routine and then the storage group ACS routine are executed.

DFDSS: For the COPY environment (&ACSENVIR = 'ALLOC') and the RESTORE environment (&ACSENVIR = 'RECOVER'), the storage class ACS routine is executed first. If the storage class is not null, the management class ACS routine and then the storage group ACS routine are executed.

For CONVERTV SMS TEST and CONVERTV SMS, (&ACSENVIR = 'CONVERT'), the storage class ACS routine is executed first. If the storage class ACS routine determines that the storage class is not null, the management class ACS routine is executed.

ISMF: When you are testing ACS routines, the data class ACS routine and then the storage class ACS routine are executed. If the storage class ACS routine determines that the storage class is not null, the management class ACS

routine and then the storage group ACS routine are executed. You can also execute each ACS routine separately from the others when performing tests.

Determining SMS Classes and Storage Groups

For each of the SMS classes, the processing is as follows:

1. If you have written an ACS routine to determine the SMS class, SMS executes the routine.

2. Next, SMS executes the corresponding **ACS installation exit**. An ACS installation exit is a program you can write to perform processing beyond the scope of the standard ACS routines. Such processing might involve:

   - Calling other programs
   - Writing SMF records
   - Writing GTF trace records
   - Performing arithmetic calculations
   - Maintaining large tables of information for quick searches
   - Taking dumps
   - Invoke ACS routines only once,

   Each exit can override the corresponding SMS class, whether explicitly specified or previously determined by an ACS routine. The exit can also invoke the corresponding ACS routine a second time, but this will not cause the installation exit to be reinvoked.

   See *MVS/ESA Data Facility Product Version 3: Customization* for additional information about ACS installation exits.

3. Finally, the system verifies that the class is defined in the currently active configuration, then performs authorization checking to verify that end users are allowed to use a selected management class or storage class.

For storage groups, SMS only invokes the ACS storage group ACS routine. Storage group does not have a corresponding ACS installation exit.

Displaying ACS Object Information

To display information about the ACS objects stored in a control data set, select option 5, DISPLAY, from the Automatic Class Selection Application Selection panel and press ENTER. You receive the ACS Object Display panel shown in Figure 43 on page 78.

```
                           ACS OBJECT DISPLAY
   COMMAND ===>

   CDS NAME: SMS.SCDS1.SCDS

   ACS RTN    SOURCE DATA SET ACS      MEMBER    LAST TRANS   LAST DATE   LAST TIME
   TYPE       ROUTINE TRANSLATED FROM  NAME      USERID       TRANSLATED  TRANSLATED
   --------   -----------------------  --------  -----------  ----------  ----------

   DATACLAS   -----------------------  --------  ---------    ----------  ----

   MGMTCLAS   AUTO.MATIC.CLASS.SELECT  --------  USERXXX      1989/01/01  12:13
              .ION.ROUTINE.MCACSRTN

   STORCLAS   DM.ACS.SOURCE            SCACSRTN  USERXXX      1989/01/01  13:13

   STORGRP    DM.ACS.SOURCE            SGACSRTN  USERXXX      1989/01/01  14:13




   USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 43. Displaying ACS Object Information

CDS NAME can be the name of the SCDS that you entered on the Automatic
Class Selection Application Selection panel or 'ACTIVE'. It identifies the data
set that contains the ACS objects that are to be displayed.

ACS RTN TYPE always lists the four types of ACS objects that can belong to an
SCDS, even if they do not actually exist in the SCDS being displayed. If they do
not exist, the corresponding display fields contain dashes.

SOURCE DATA SET ACS ROUTINE TRANSLATED FROM contains the name of
the data set that has the source code responsible for creating the ACS object.
As shown in the figure, the data set name is folded if it is more than 23 charac-
ters in length. If an ACS object of a particular type does not exist in the SCDS
being displayed, then this field contains dashes.

MEMBER NAME displays the member within the data set that contains the ACS
source code used to create the corresponding ACS object. If an ACS object of
the particular type does not exist in the SCDS being displayed, or if the
SOURCE DATA SET ACS ROUTINE TRANSLATED FROM is not a PDS, the field
contains dashes.

LAST TRANS USERID displays the TSO userid of the person who last translated
the ACS routine. If an ACS object of the particular type does not exist, the field
contains blanks.

LAST DATE TRANSLATED displays the date when the corresponding ACS object
was created. If an ACS object of the particular type does not exist, the field
contains blanks.

LAST TIME TRANSLATED displays the time when the corresponding ACS object
was created. If an ACS object of the particular type does not exist, the field
contains blanks.

Deleting an ACS Object from an SCDS

To delete an ACS object from an SCDS, select option 6, DELETE, from the Automatic Class Selection Application Selection panel and press ENTER. You receive the Delete ACS Object panel shown in Figure 44.

```
                         DELETE ACS OBJECT
  COMMAND ===>

  TO DELETE AN ACS OBJECT, SPECIFY:

     SCDS NAME        ===> 'SMS.SCDS1.SCDS'
                               (1 to 44 character data set name)


     ACS ROUTINE TYPE ===> SG    (DC=Data Class, MC=Management Class,
                                  SC=Storage Class, SG=Storage Group)







  USE ENTER TO PERFORM ACS DELETION;
  USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 44. Deleting ACS Objects

SCDS NAME contains the name of the SCDS from which the ACS object is to be deleted. It is a required field and is primed with the last referenced SCDS. The default is blanks.

ACS ROUTINE TYPE contains the type of ACS object. It is a required field and it is primed with the last used value. The default is blanks.

When you press ENTER, you receive the CONFIRM DELETE REQUEST panel to certify that you do indeed want to delete the ACS object.
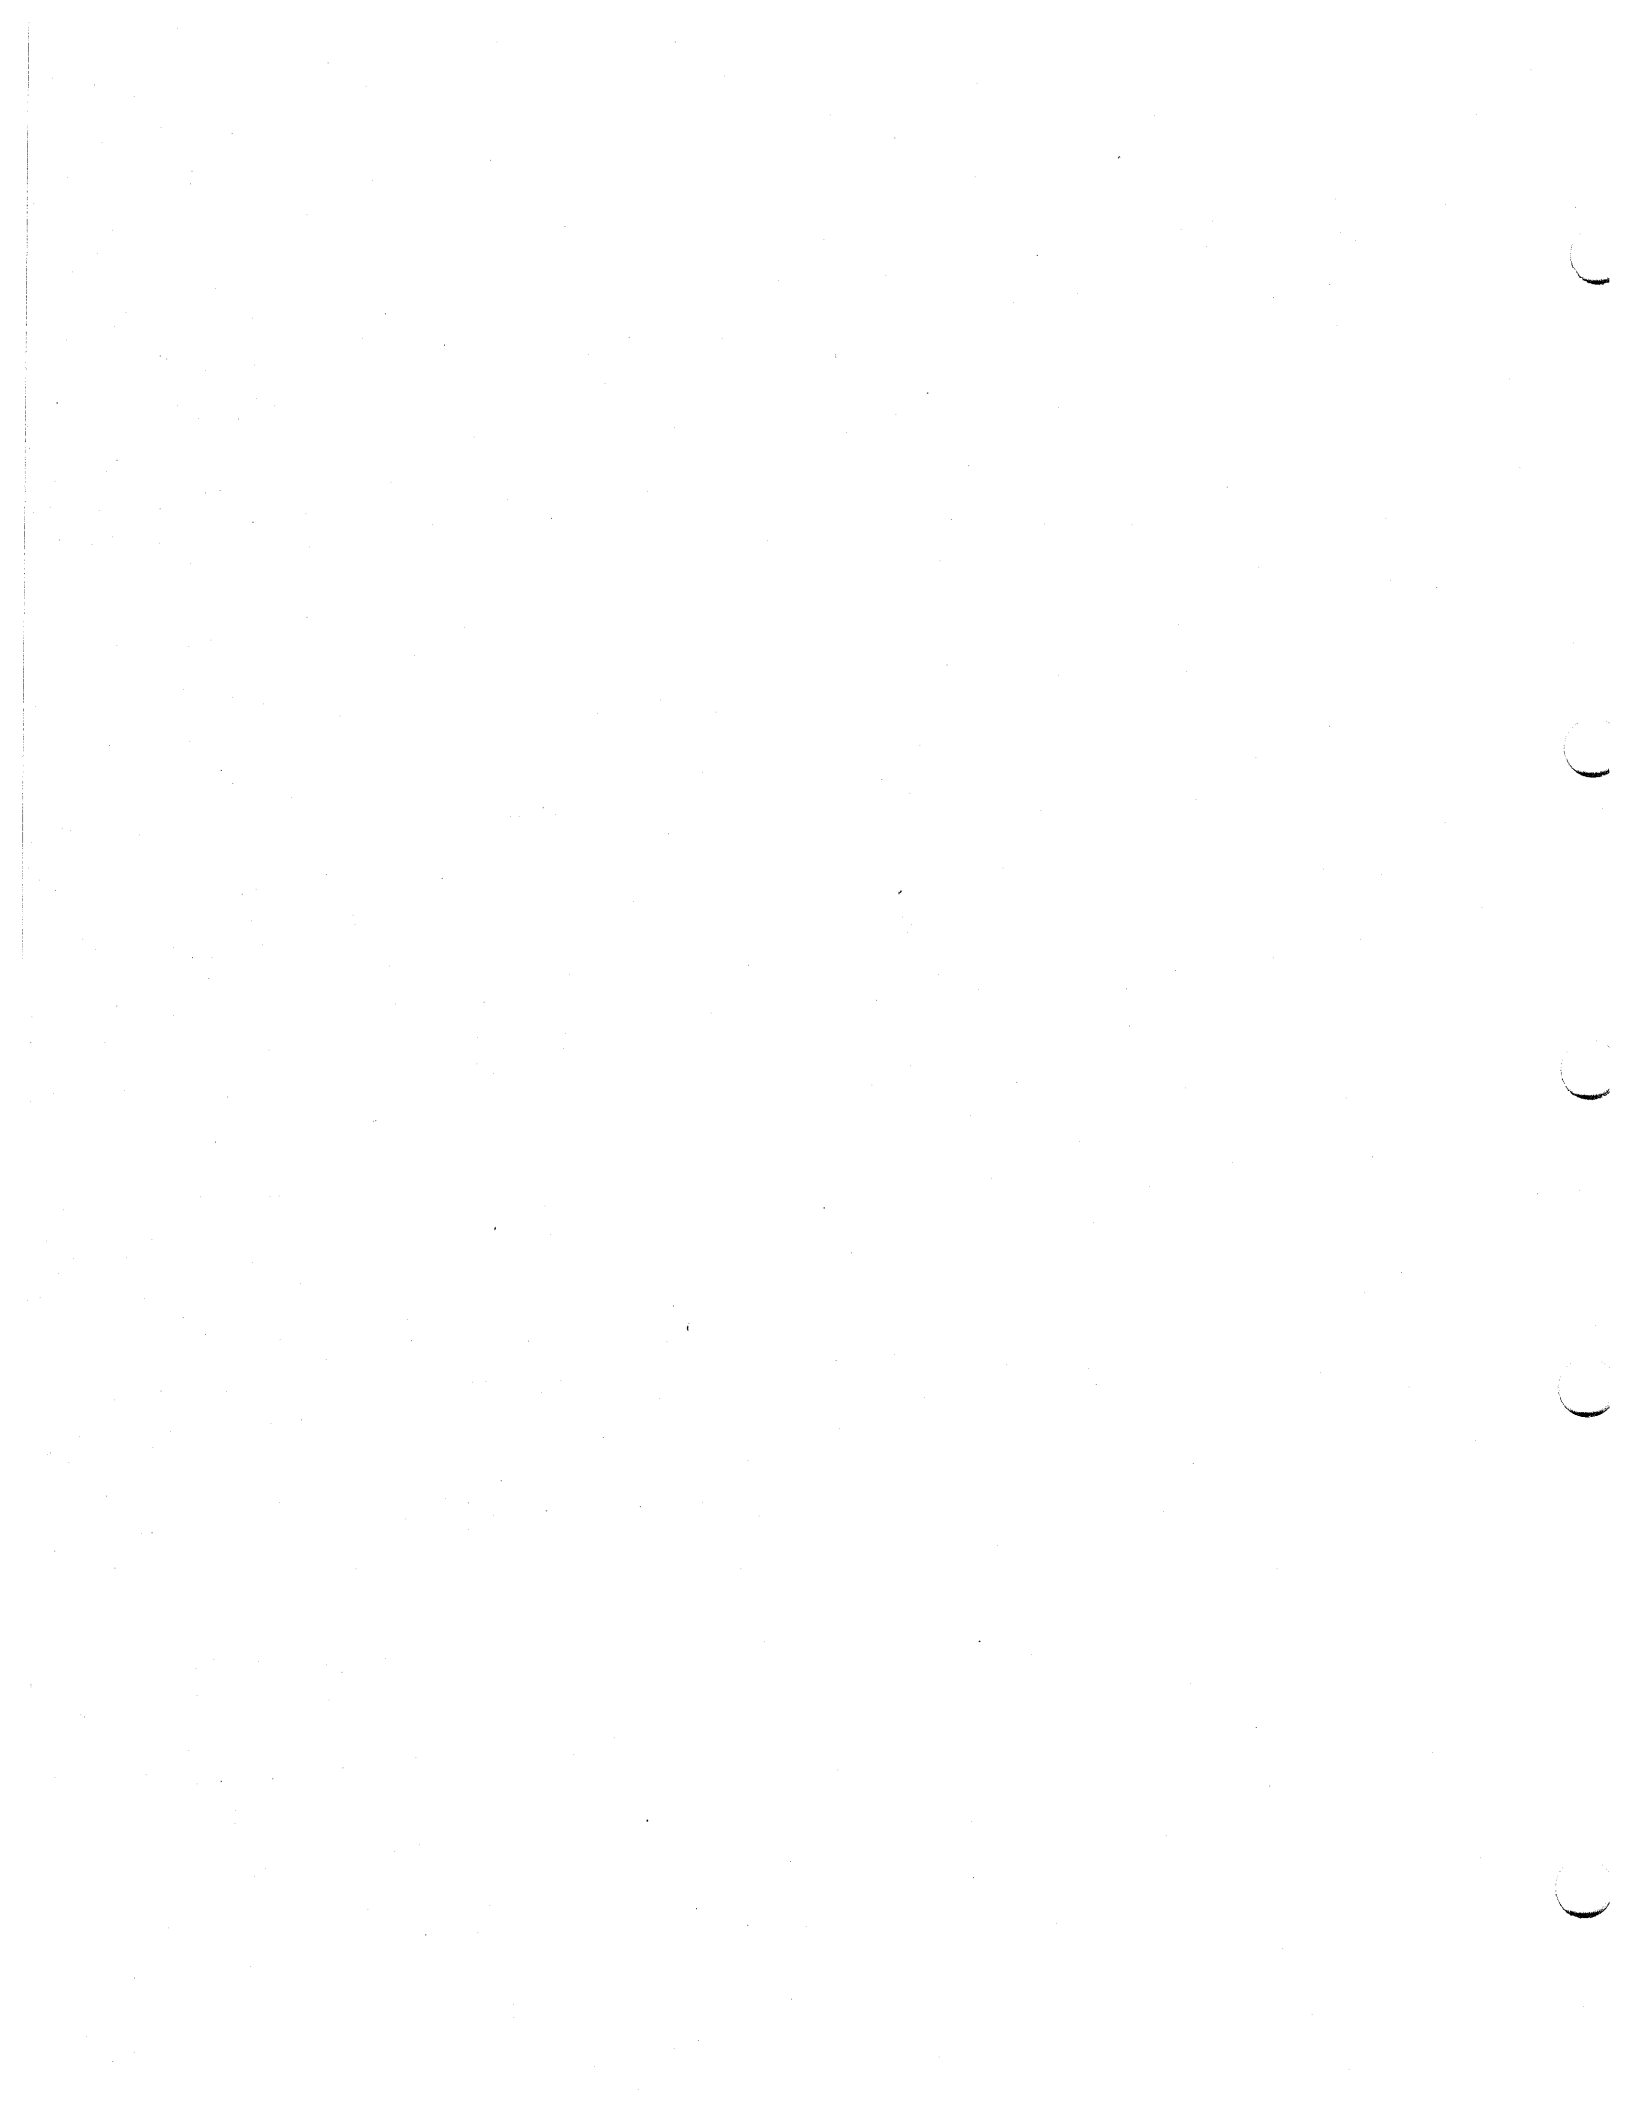
Chapter 9. Activating SMS Configurations

You can manually activate an SMS configuration, or you can automatically activate it at IPL. The first part of this chapter explains how to perform the initial activation of an SMS configuration using a four-step manual approach. The second part of this chapter explains how to activate an SMS configuration automatically at future IPLs. The final part of this chapter explains how to change individual SMS parameters using the SETSMS operator command.

Manually Activating the First SMS Configuration

IGDSSIIN is the subsystem initialization routine module for SMS. By omitting it from the SMS record of IEFSSNnn for each system in the SMS complex, you can manually control the activation of an SMS configuration. Refer to Chapter 2, "Preparing for SMS" on page 7.

Step One—IPL Each System in the SMS Complex

After defining SMS as a subsystem to MVS, IPL each system in the SMS complex. The presence of the SMS record in IEFSSNnn tells MVS to recognize SMS as a valid subsystem within each system. The absence of the IGDSSIIN module name in IEFSSNnn tells the system that you want to manually activate a configuration.

Step Two—Prepare One System

From one system in the SMS complex, issue the T SMS = mm command. 'mm' identifies IGDSMSmm as the SMS initialization control member of SYS1.PARMLIB. SMS uses the ACDS and COMMDS identified in IGDSMSmm to manage storage. Because the initial ACDS (and COMMDS) are currently empty, the system is activated with a **null configuration**. With a null configuration, this system continues to function as a non-SMS system.

Step Three—Activate a Configuration for One System

After activating SMS with a null configuration, activate an SMS configuration contained in a valid SCDS on the same system. You can use either the ISMF ACTIVATE command or the SETSMS operator command. Both procedures copy the contents of the SCDS to the ACDS specified in IGDSMSmm.

Activating with the ISMF ACTIVATE Command: On the Control Data Application Selection panel shown in Figure 6 on page 19, specify the name of an SCDS and issue the ACTIVATE command from the command line. A *Write to Programmer* message indicates if the activation is successful, provided you have WTPMSG in the TSO/E PROFILE.

Activating with the SETSMS Operator Command: From the operator console, issue the command:

```
SETSMS SCDS(dsname)
```

where dsname identifies the name of the SCDS to be activated. Using the SCDS defined earlier in this manual, the command would be SETSMS SCDS(SMS.SCDS1.SCDS).

Step Four—Activate the Configuration on the Other Systems

For the other systems in the SMS complex, use the T SMS = mm command to activate the SMS configuration identified in Step Three. In each system, IGDSMSmm specifies the name of the ACDS containing the SMS configuration. All of the IGDSMSmm members must point to the same ACDS. Because the ACDS is no longer empty, the systems use it (and the COMMDS) to manage storage.

Automatically Activating an SMS Configuration

After you have activated an SMS configuration and thereby copied it into an ACDS, you can automatically activate the ACDS at future IPLs. For each system in the SMS complex, update the SMS record in IEFSSNnn to include the IGDSSIIN module name.

Make certain that each ID field identifies the IGDSMSmm member containing the name of the last used ACDS. All of the IGDSMSmm members must point to the same ACDS. At future IPLs, the SMS configuration contained in the ACDS is activated by all systems in the SMS complex. For example:

```
SMS,IGDSSIIN,'ID=02,PROMPT=DISPLAY'
```

indicates that the ACDS specified in IGDSMS02 contains the SMS configuration to be activated at future IPLs.

Changing SMS Parameters

Once you have activated an SMS configuration, you can use the SETSMS operator command to change SMS parameters. The SETSMS operator command has the following options:

SCDS(*dsname*)
> *dsname* identifies the name of the data set that SMS is to activate. If the SCDS is invalid, the command will fail.

ACDS(*dsname*)
> *dsname* identifies the name of the SCDS data set containing the active configuration. If *dsname* contains an invalid configuration, then you must also specify an SCDS using the SCDS option. If both the SCDS and the ACDS *dsname*s are invalid, the command will fail.

SAVEACDS(*dsname*)
> *dsname* identifies the name of the data set to contain a copy of the current active configuration.

COMMDS(*dsname*)
> *dsname* represents the name of the COMMDS.

INTERVAL(*nnn*)
> *nnn* represents the **synchronization time interval** of the system, which is the number of seconds that SMS allows before this system checks the COMMDS for news from other systems in the SMS complex. You can specify a value from 1 to 999.

DINTERVAL(*nnn*)
> *nnn* represents the number of seconds SMS allows to elapse before it reads device statistics. You can specify a value from 1 to 999.

TRACE(OFF|ON)

The **TRACE** keyword indicates whether you want SMS to perform tracing.

SIZE(*nnn*|*nnn*K|*nnn*M)

The **SIZE** keyword specifies the size of the trace table in bytes. When the unit is **K**, values can range from 0K to 255,000K, and they are rounded up to the nearest 4K unit. When the unit is **M**, values can range from 0M to 255M. If you specify a value of 0, no tracing is performed.

TYPE(ALL|ERROR)

The **TYPE** keyword specifies whether you want to trace all events or only errors.

JOBNAME(*jobname*|*)

The **JOBNAME** keyword limits tracing to a certain job or permits tracing on all jobs.

ASID(*asid*|*)

The **ASID** keyword limits tracing to a certain address space or permits it for all address spaces.

SELECT(*event, event, ...*|ALL)

The **SELECT** keyword adds items to the list of events to be traced.

DESELECT(*event, event, ...*)

The **DESELECT** keyword deletes items from the list of events to be traced. If you specify SELECT and DESELECT keyword events that conflict, the keyword that appears last has final authority.

If you do not specify a keyword value, SMS uses the current value (specified in IGDSMSmm or by a previous operator command). For a list of the events that you can specify on the SELECT and DESELECT keywords, see "Initializing SMS through IGDSMSmm" on page 11.

Considerations when Changing SMS Configurations

When activating a new SMS configuration, you have two alternatives for keeping the currently active SMS configuration information. One alternative is to keep, but never modify, the original SCDS from which the current SMS configuration was activated. You also need to maintain a log of all status changes (VARY storage group, ENABLE, etc.) you make to the currently active SMS configuration. If in the future you activate a different SMS configuration but then decide you want to fall back to your original, you can re-activate the SCDS. You lose all the status changes you have made since activating the SCDS, and you must now re-enter them. But you return to the original SMS configuration.

The better alternative is to save the current inactive SMS configuration using the SETSMS operator command:

```
SETSMS SAVEACDS(ACDS.FALLBACK)
```

This not only saves the current storage management policy in
ACDS.FALLBACK, but it also saves the status changes you have made since
the original SCDS was activated. You can then activate the new SMS config-
uration. If in the future you decide you want to fall back to the original SMS
configuration, you can use the SETSMS operator command:

```
SETSMS ACDS(ACDS.FALLBACK)
```

to re-activate it. This alternative is also useful if you have altered the SCDS
that you originally activated.

Chapter 10. Maintaining SMS

After building a configuration and activating SMS, you need to monitor and adjust it over time. This chapter explains how to maintain SMS after you have activated it.

Displaying SMS Information

You can use the DISPLAY operator command to gather SMS information. You can also enter the LISTSYS and LISTVOL line operators on the ISMF Storage Group List panels to get information about the currently active configuration, storage groups, and storage volumes.

Displaying Information about the Active Configuration

To display information about the currently active configuration, issue the DISPLAY command:

D SMS,ACTIVE or D SMS

where D is an abbreviation for DISPLAY. An example of the generated output appears in Figure 45.

```
09.23.36            IGD002I 09:23:36 DISPLAY SMS 401

SCDS = SMS.SCDS1.SCDS
ACDS = SMS.ACDS1.ACDS
COMMDS = SMS.COMMDS1.COMMDS

DINTERVAL = 150
REVERIFY = NO
ACSDEFAULTS = NO
     SYSTEM      CONFIGURATION LEVEL      INTERVAL SECONDS
     SYSTEM01    1987/12/01 09:23:29           15
     SYSTEM02    ---------- --------          N/A
     SYSTEM03    ---------- --------          N/A
     SYSTEM04    ---------- --------          N/A
     SYSTEM05    ---------- --------          N/A
     SYSTEM06    ---------- --------          N/A
     SYSTEM07    ---------- --------          N/A
     SYSTEM08    ---------- --------          N/A
```

Figure 45. Displaying Information about the Active Configuration

The heading of this output indicates the time (09.23.36), the command issued (DISPLAY SMS), and a three digit decimal identification associated with the display (401).

The next section indicates the names of the current control data sets. The naming convention used here specifies a first level qualifier of SMS for all SMS control data sets. The second level qualifier identifies the type of SMS data set. The third level qualifier uniquely identifies the data set.

The last portion of the output indicates configuration levels, relative to the system from which you issued the DISPLAY command. This last portion also contains the synchronization time interval for each system.

Displaying the SMS TRACE Command

To display information about the SMS trace options, issue the DISPLAY command:

```
D SMS,TRACE
```

where D is an abbreviation for DISPLAY. An example of the generated output appears in Figure 46.

```
 09.25.43   DISPLAY SMS 401

 TRACE    = {ON|OFF}   SIZE =nnnK      TYPE {ERROR|ALL}
    JOBNAME ={jjj|*}    ASID ={asid|*}
    TRACING EVENTS:
 MODULE ={ON|OFF}  SMSSJF ={ON|OFF}  SMSSI  ={ON|OFF}  ACSINT ={ON|OFF}
 OPCMD  ={ON|OFF}  CONFC  ={ON|OFF}  CDSC   ={ON|OFF}  CONFS  ={ON|OFF}
 MSG    ={ON|OFF}  ERR    ={ON|OFF}  CONFR  ={ON|OFF}  CONFA  ={ON|OFF}
 ACSPRO ={ON|OFF}  IDAX   ={ON|OFF}  DISP   ={ON|OFF}  CATG   ={ON|OFF}
 VOLREF ={ON|OFF}  SCHEDP ={ON|OFF}  SCHEDS ={ON|OFF}  VTOCL  ={ON|OFF}
 VTOCD  ={ON|OFF}  VTOCR  ={ON|OFF}  VTOCC  ={ON|OFF}  VTOCA  ={ON|OFF}
 RCD    ={ON|OFF}
```

Figure 46. Displaying Trace Information

The heading of this output indicates the time (09.25.43), the command was issued (DISPLAY SMS), and a three digit decimal identification associated with the display (401).

The next section indicates the status of the SMS trace option, the size of the SMS trace table, and the type of SMS trace entries. ERROR means only error type trace entries are traced; ALL indicates all types of trace entries are traced. This section also includes the jobname which indicates the tracing scope in relation to jobs being run, and ASID which indicates the tracing scope in relation to address spaces. ASID means tracing is limited to a particular address space and * indicates tracing is performed for all address spaces.

The last section, TRACING EVENTS, indicates which SMS events are selected for tracing.

You can use this tracing display when you have been requested to collect information by the IBM Support Center for diagnosing problems. See the *DFP: Diagnosis Reference* for additional information.

Displaying Storage Group Status Using the DISPLAY Command

To display information about the status of a storage group, issue the following command:

```
D SMS,STORGRP(storage-group-name)
```

For storage-group-name you specify the name of a single storage group or you specify ALL to display the status of all storage groups. An example of the generated output for the single storage group SG1 appears in Figure 47 on page 87.

```
09.26.17          IGD002I 09:26:17 DISPLAY SMS 407
STORGRP           SYSTEM= 1 2 3 4 5 6 7 8
SG1                      + . . . . . . .
**************************** LEGEND ****************************
. THE STORAGE GROUP OR VOLUME IS NOT DEFINED TO THE SYSTEM
+ THE STORAGE GROUP OR VOLUME IS ENABLED
- THE STORAGE GROUP OR VOLUME IS DISABLED
* THE STORAGE GROUP OR VOLUME IS QUIESCED
D THE STORAGE GROUP OR VOLUME IS DISABLED FOR NEW ALLOCATIONS ONLY
Q THE STORAGE GROUP OR VOLUME IS QUIESCED FOR NEW ALLOCATIONS ONLY
SYSTEM 1 = SYSTEM01     SYSTEM 5 = SYSTEM05
SYSTEM 2 = SYSTEM02     SYSTEM 6 = SYSTEM06
SYSTEM 3 = SYSTEM03     SYSTEM 7 = SYSTEM07
SYSTEM 4 = SYSTEM04     SYSTEM 8 = SYSTEM08
```

Figure 47. Displaying Storage Group Status Information

The heading of this output indicates the time and the command issued. Next, the output lists the status of the storage group.

Specifying the optional LISTVOL parameter provides the status of the volumes associated with the storage group. If you have volume 123456 in storage group SG1, the command:

D SMS,STORGRP(SG1),LISTVOL

generates output similar to that shown in Figure 48.

```
09.28.38          IGD002I 09:28:38 DISPLAY SMS 411
STORGRP           SYSTEM= 1 2 3 4 5 6 7 8
SG1                      + . . . . . . .
VOLUME    UNIT   SYSTEM= 1 2 3 4 5 6 7 8 STORGRP NAME
123456                   + . . . . . . .    SG1
**************************** LEGEND ****************************
. THE STORAGE GROUP OR VOLUME IS NOT DEFINED TO THE SYSTEM
+ THE STORAGE GROUP OR VOLUME IS ENABLED
- THE STORAGE GROUP OR VOLUME IS DISABLED
* THE STORAGE GROUP OR VOLUME IS QUIESCED
D THE STORAGE GROUP OR VOLUME IS DISABLED FOR NEW ALLOCATIONS ONLY
Q THE STORAGE GROUP OR VOLUME IS QUIESCED FOR NEW ALLOCATIONS ONLY
SYSTEM 1 = SYSTEM01     SYSTEM 5 = SYSTEM05
SYSTEM 2 = SYSTEM02     SYSTEM 6 = SYSTEM06
SYSTEM 3 = SYSTEM03     SYSTEM 7 = SYSTEM07
SYSTEM 4 = SYSTEM04     SYSTEM 8 = SYSTEM08
```

Figure 48. Displaying Storage Group Volume Status Information

Displaying Storage Group Status Using ISMF

You can display a storage group's status with respect to each system in the SMS complex using the LISTSYS line operator in ISMF. The storage group must be a pool type, and it must be part of the active configuration. Go to the Storage Group Application Selection panel, shown in Figure 8 on page 26, from the ISMF Primary Option Menu. Then create a list of storage groups. For example, specifying the following information on the Storage Group Application Selection panel will create a list of storage groups whose names begin with SGNAME:

```
CDS NAME            ===> 'ACTIVE'
STORAGE GROUP NAME  ===> SGNAME*
SELECT OPTION       ===> 1
```

Press the ENTER key to display the Storage Group List panel, and enter LISTSYS in the Line Operator column next to the desired pool type storage group. Press ENTER to display the Storage Group System Status Display panel shown in Figure 49.

Note: If a volume has been initialized, the storage group must be defined in the active configuration in order for space information to be displayed when the LISTSYS line operator is issued.

```
                   SMS STORAGE GROUP STATUS DISPLAY
COMMAND ===>

CDS NAME:           ACTIVE
STORAGE GROUP NAME: SGNAME01
TOTAL SPACE FOR STORAGE GROUP
         MB-TOTAL:   200000
         MB-FREE:    100000
         % FREE:         50

                       ------ GROUP SPACE ---------
     SYSTEM   SMS SG    AVAILABLE FOR ALLOCATION
     NAME     STATUS   MB-TOTAL   MB-FREE   % FREE
   --------  -------   --------- ---------  ------
     SYSTEM01 ENABLE    100000     50000       50
     SYSTEM02 ENABLE    100000     50000       50
     SYSTEM03 NOTCON    ------     -----       --
     SYSTEM04 NOTCON    ------     -----       --
     SYSTEM05 NOTCON    ------     -----       --
     SYSTEM06 NOTCON    ------     -----       --
     SYSTEM07 NOTCON    ------     -----       --
     SYSTEM08 NOTCON    ------     -----       --


  USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 49. Displaying Storage Group Status for each System

TOTAL SPACE FOR STORAGE GROUP is the total capacity of all the volumes belonging to storage group SGNAME01, if all the volumes have a status of ENABLE.

MB-TOTAL is the total number of megabytes belonging to the storage group.

MB-FREE is the total number of free megabytes belonging to the storage group.

% FREE is the percentage of free storage group space.

GROUP SPACE AVAILABLE FOR ALLOCATION represents the sum of all volume space that is currently available to each system online to MVS, volume is evoked, and the storage group is enabled. and enabled to the storage group. The fields contain values only if the SMS SG STATUS is ENABLE.

MB-TOTAL is the total available space in megabytes.

MB-FREE is the total available space in megabytes that is free.

% FREE is the percentage of total available space that is free.

Displaying Volumes Using the DISPLAY Command

To display information about the status of an individual volume, issue the following command:

D SMS,VOLUME(serial-number)

An example of the generated output for volume serial number 123456 appears on the operator display panel shown in Figure 50.

```
      09.30.23        IGD002I 09:30:23 DISPLAY SMS 415
   00 VOLUME    UNIT    SYSTEM= 1 2 3 4 5 6 7 8 STORGRP NAME
      123456                    + . . . . . . .   SG1
      ***************************** LEGEND *****************************
      . THE STORAGE GROUP OR VOLUME IS NOT DEFINED TO THE SYSTEM
      + THE STORAGE GROUP OR VOLUME IS ENABLED
      - THE STORAGE GROUP OR VOLUME IS DISABLED
      * THE STORAGE GROUP OR VOLUME IS QUIESCED
      D THE STORAGE GROUP OR VOLUME IS DISABLED FOR NEW ALLOCATIONS ONLY
      Q THE STORAGE GROUP OR VOLUME IS QUIESCED FOR NEW ALLOCATIONS ONLY
      SYSTEM 1 = SYSTEM01    SYSTEM 5 = SYSTEM05
      SYSTEM 2 = SYSTEM02    SYSTEM 6 = SYSTEM06
      SYSTEM 3 = SYSTEM03    SYSTEM 7 = SYSTEM07
      SYSTEM 4 = SYSTEM04    SYSTEM 8 = SYSTEM08
```

Figure 50. Displaying Status Information for Individual Volumes

Displaying the Status of a Device

You can use the DEVSERV command to display the status of a device. The command has the syntax shown in Figure 51.

```
[DEVSERV]  [SMS]  ,  uuu  [,nn]  [,ONLINE ]  [,L= [a  ] ]
[DS     ]  [S  ]               [,ON     ]      [   [cc ] ]
                               [,OFFLINE]      [   [cca] ]
                               [,OFF    ]
```

Figure 51. DEVSERV Command Syntax

uuu   The target device number for which information is requested.

nn   An optional decimal number from 1 to 32 that specifies the
              number of devices, beginning with the target device, for which
              information is requested. The default is 1.

ONLINE/ON   Restricts the inquiry to online devices only.

OFFLINE/OFF Restricts the inquiry to offline devices only.

L=a|cc|cca Specifies where the results of the inquiry are to be displayed:
              the display area (a), the console (cc), or both (cca).

The command DEVSERV S,430 produces the output shown in Figure 52 for target
device 430.

```
16.24.26          devserv s,430
16.24.26          IGD001I 16:24:26 DEVSERV SMS 569
UNIT,TYPE      ,M,VOLUME,VOLSTAT      STORGRP,SGSTAT
430 ,3380      ,0,XP0101,ENABLED      SXP01,QUIESCED

*********************** SYMBOL DEFINITIONS ***********************
0 = ONLINE                          + = PATH AVAILABLE
```

Figure 52. Displaying the Status of a Device

Listing SMS Classes and Storage Groups

You can partially specify SMS class or storage group names to generate a
listing of related items. You can then enter the following commands for any
item on the list.

Figure 53. LIST Commands				
CAncel	CLEar	DOwn	END	ERb
FINd	HELP	Left	Profile	RESHow
RETURN	RIght	SAve	SOrt	Up

In Figure 53, uppercase letters indicate the minimum allowable abbreviation.
Appendix A, "ISMF Command and Line Operator Reference Summary" on
page 139 summarizes all the ISMF commands and line operators, describes
them, and lists the applications from which you can issue them.

Listing Storage Groups

If you choose the List option from the Storage Group Application Selection
panel, shown in Figure 8 on page 26, you can select a sorting order by
entering Y, yes, in the RESPECIFY SORT CRITERIA field. ISMF displays the
Storage Group Sort Entry Panel shown in Figure 54 on page 91.

```
                        STORAGE GROUP SORT ENTRY PANEL
COMMAND ===>

SPECIFY ONE OR MORE ATTRIBUTE NUMBERS FOR SORT SEQUENCE:
   MAJOR FIELD ===> 2     MINOR FIELD 1 ===>        MINOR FIELD 2 ===>


SPECIFY A FOR ASCENDING OR D FOR DESCENDING SORT ORDER:
   MAJOR FIELD ===> A     MINOR FIELD 1 ===>        MINOR FIELD 2 ===>

     (1)  LINE OPERATOR              (13)  MIGRATION LOW
     (2)  STORAGE GROUP NAME         (14)  MB TOTAL SPACE
     (3)  STORAGE GROUP TYPE         (15)  MB FREE SPACE
     (4)  VIO MAXSIZE                (16)  % FREE SPACE
     (5)  VIO UNIT                   (17)  LAST MODIFIED USERID
     (6)  AUTO MIGRATE               (18)  LAST DATE MODIFIED
     (7)  MIGRATE SYSTEM             (19)  LAST TIME MODIFIED
     (8)  AUTO BACKUP                (20)  DUMP CLASS
     (9)  BACKUP SYSTEM              (21)  DUMP CLASS
    (10)  AUTO DUMP                  (22)  DUMP CLASS
    (11)  DUMP SYSTEM                (23)  DUMP CLASS
    (12)  MIGRATION HIGH             (24)  DUMP CLASS


USE ENTER TO PERFORM SORT;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 54. Specifying Storage Group Sort Criteria

As shown in Figure 54, the default sort order is based on Storage Group Name,
2, in ascending order, A. If you specify N, no, in the RESPECIFY SORT CRI-
TERIA field, ISMF uses the most recent sorting order. Based on the sorting cri-
teria, ISMF displays a list of storage groups like the one shown in Figure 55.
On the listing, you can enter line operators such as COPY, ALTER, DELETE, and
HIDE against individual entries.

```
                          STORAGE GROUP LIST          SORT WAS SUCCESSFUL
COMMAND ===>                                           SCROLL ===> PAGE
                                                      Entries 1-1 of 1
                                                      Data Columns 3-9 of 24

CDS NAME:  SMS.SCDS1.SCDS

ENTER LINE OPERATORS BELOW:

      LINE      STORGRP   SG     VIO      VIO  AUTO  MIGRATE   AUTO  BACKUP
      OPERATOR  NAME      TYPE   MAXSIZE  UNIT MIGR  SYSTEM    BACK  SYSTEM
      ---(1)----  --(2)---  -(3)-  --(4)--  (5)- (6)-  --(7)---  (8)- --(9)---
                PRIMARY   POOL   -------  ---- YES   --------  YES  --------
      ----------  --------  ------ BOTTOM OF  DATA ------  --------  ----------
```

Figure 55. Displaying a List of Storage Groups

You can scroll up, down, left, or right. However, ISMF always displays the LINE
OPERATOR and STORGRP NAME fields on the left side of the panel.

Listing Management Classes

If you choose the List option from the Management Class Application Selection panel, shown in Figure 15 on page 38, you can select a sorting order by specifying Y, yes, in the RESPECIFY SORT CRITERIA field. ISMF displays the Management Class Sort Entry Panel in Figure 56.

```
                     MANAGEMENT CLASS SORT ENTRY PANEL
COMMAND ===>

SPECIFY ONE OR MORE ATTRIBUTE NUMBERS FOR SORT SEQUENCE:
   MAJOR FIELD ===> 2     MINOR FIELD 1 ===>        MINOR FIELD 2 ===>

SPECIFY A FOR ASCENDING OR D FOR DESCENDING SORT ORDER:
   MAJOR FIELD ===> A     MINOR FIELD 1 ===>        MINOR FIELD 2 ===>

   (1)  LINE OPERATOR               (12) BACKUP FREQUENCY
   (2)  MANAGEMENT CLASS NAME       (13) NUMBER OF BACKUPS (DATA SET EXISTS)
   (3)  EXPIRE AFTER DAYS NON-USAGE (14) NUMBER OF BACKUPS (DATA SET DELETED)
   (4)  EXPIRE AFTER DATE/DAYS      (15) RETAIN DAYS ONLY BACKUP
   (5)  RETENTION LIMIT             (16) RETAIN DAYS EXTRA BACKUPS
   (6)  PARTIAL RELEASE             (17) ADMIN OR USER COMMAND BACKUP
   (7)  PRIMARY DAYS NON-USAGE      (18) AUTO BACKUP
   (8)  LEVEL ONE DAYS NON-USAGE    (19) LAST MODIFIED USERID
   (9)  COMMAND OR AUTO MIGRATE     (20) LAST DATE MODIFIED
   (10) # GDG ELEMENTS ON PRIMARY   (21) LAST TIME MODIFIED
   (11) ROLLED-OFF GDS ACTION


USE ENTER TO PERFORM SORT;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 56. Specifying Management Class Sort Criteria

You need to specify either an ascending or descending order sequence for each of the sorting fields. ISMF saves the sorting criteria for the life of the application, or until you update the criteria. If you specify N, no, in the RESPECIFY SORT CRITERIA field, ISMF uses the most recent sorting order.

Based on the sorting criteria, ISMF displays a list of management classes like the one shown in Figure 57 on page 93. On the listing, you can enter line operators such as COPY, DISPLAY, and HIDE against individual entries. But some line operators cannot be used with ACTIVE such as ALTER and DELETE.

```
┌─────────────────────────────────────────────────────────────────────────────┐
│                         MANAGEMENT CLASS LIST           SORT WAS SUCCESSFUL   │
│   COMMAND ===>                                              SCROLL ===> HALF   │
│                                                         Entries 1-8 of 8       │
│                                                         Data Columns 3-7 of 21 │
│   CDS NAME:  SMS.SCDS1.SCDS                                                    │
│                                                                                │
│   ENTER LINE OPERATORS BELOW:                                                 │
│                                                                                │
│       LINE       MGMTCLAS  EXPIRE     EXPIRE        RET      PARTIAL  PRIMARY  │
│       OPERATOR   NAME      NON-USAGE  DATE/DAYS     LIMIT    RELEASE  DAYS     │
│       ---(1)---- --(2)---  ---(3)---  ---(4)----   --(5)--  --(6)--  ---(7)-- │
│                  TATATA    NOLIMIT       NOLIMIT   NOLIMIT  NO             2   │
│                  TEST4     NOLIMIT       NOLIMIT      9999  NO             2   │
│                  TEST5     NOLIMIT       NOLIMIT   NOLIMIT  NO             2   │
│                  T0        NOLIMIT       NOLIMIT   NOLIMIT  NO             2   │
│                  T1        NOLIMIT    1986/12/31   NOLIMIT  NO             2   │
│                  T2        NOLIMIT       NOLIMIT   NOLIMIT  NO             2   │
│                  T3        NOLIMIT       NOLIMIT   NOLIMIT  NO            14   │
│                  T9        NOLIMIT       NOLIMIT   NOLIMIT  NO             2   │
│       ---------- --------  ------ BOTTOM OF DATA ------ -------- ---------- │
│                                                                                │
│                                                                                │
│   USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.                          │
└─────────────────────────────────────────────────────────────────────────────┘
```

Figure 57. Displaying a List of Management Classes

You can scroll up, down, left, and right. However, ISMF always displays the
LINE OPERATOR and MGMTCLAS fields on the left side of the panel.

Listing Storage Classes

If you choose the List option from the Storage Class Application Selection
panel, shown in Figure 20 on page 48, you can select a sorting order by speci-
fying Y, yes, in the RESPECIFY SORT CRITERIA field. ISMF displays the Storage
Class Sort Entry Panel in Figure 58.

```
┌─────────────────────────────────────────────────────────────────────────────┐
│                         STORAGE CLASS SORT ENTRY PANEL                         │
│   COMMAND ===>                                                                 │
│                                                                                │
│   SPECIFY ONE OR MORE ATTRIBUTE NUMBERS FOR SORT SEQUENCE:                     │
│     MAJOR FIELD ===> 2      MINOR FIELD 1 ===>        MINOR FIELD 2 ===>       │
│                                                                                │
│   SPECIFY A FOR ASCENDING OR D FOR DESCENDING SORT ORDER:                      │
│     MAJOR FIELD ===> A      MINOR FIELD 1 ===>        MINOR FIELD 2 ===>       │
│                                                                                │
│      (1)    LINE OPERATOR                                                      │
│      (2)    STORAGE CLASS NAME                                                 │
│      (3)    DIRECT MILLISECOND RESPONSE                                        │
│      (4)    DIRECT BIAS                                                        │
│      (5)    SEQUENTIAL MILLISECOND RESPONSE                                    │
│      (6)    SEQUENTIAL BIAS                                                    │
│      (7)    AVAILABILITY                                                       │
│      (8)    GUARANTEED SPACE                                                   │
│      (9)    LAST MODIFIED USERID                                               │
│     (10)    LAST DATE MODIFIED                                                 │
│     (11)    LAST TIME MODIFIED                                                 │
│                                                                                │
│                                                                                │
│   USE ENTER TO PERFORM SORT;                                                   │
│   USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.                          │
└─────────────────────────────────────────────────────────────────────────────┘
```

Figure 58. Specifying Storage Class Sort Criteria

You need to specify either an ascending or descending order sequence for each of the sorting fields. ISMF saves the sorting criteria for the life of the application, or until you update the criteria. If you specify N, no, in the RESPECIFY SORT CRITERIA field, ISMF uses the most recent sorting order.

Based on the sorting criteria, ISMF displays a list of storage classes like the one shown in Figure 59. On the listing, you can enter line operators such as COPY, ALTER, DELETE, DISPLAY, and HIDE against individual entries.

```
                              STORAGE CLASS LIST
   COMMAND ===>                                          SCROLL ===> HALF
                                                     Entries 1-1 of 1
                                                     Data Columns 3-8 of 11

   CDS NAME:  SMS.SCDS1.SCDS

   ENTER LINE OPERATORS BELOW:

     LINE        STORCLAS  DIR RESP  DIR    SEQ RESP  SEQ    AVAIL-      GUARANTEED
     OPERATOR    NAME      (MSEC)    BIAS   (MSEC)    BIAS   ABILITY     SPACE
     ---(1)----  --(2)---  --(3)---  (4)-   --(5)---  (6)-   ---(7)----  ---(8)----
                 DBFAST          5   R             5   R    STANDARD    YES
     ----------  --------  ------ BOTTOM  OF   DATA  ------  --------    ----------




   USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 59. Displaying a List of Storage Classes

You can scroll up, down, left, and right. However, ISMF always displays the LINE OPERATOR and STORCLAS NAME fields on the left side of the panel.

Listing Data Classes

If you choose the List option from the Data Class Application Selection panel, shown in Figure 22 on page 54, you can select a sorting order by specifying Y, yes, in the RESPECIFY SORT CRITERIA field. ISMF displays the Data Class Sort Entry Panel in Figure 60 on page 95.

```
                          DATA CLASS SORT ENTRY PANEL
COMMAND ===>


SPECIFY ONE OR MORE ATTRIBUTE NUMBERS FOR SORT SEQUENCE:
   MAJOR FIELD ===> 2      MINOR FIELD 1 ===>         MINOR FIELD 2 ===>


SPECIFY A FOR ASCENDING OR D FOR DESCENDING SORT ORDER:
   MAJOR FIELD ===> A      MINOR FIELD 1 ===>         MINOR FIELD 2 ===>

   (1)   LINE OPERATOR            (13)  RETPD OR EXPDT
   (2)   DATA CLASS NAME          (14)  VOLUME COUNT
   (3)   RECORD ORGANIZATION      (15)  IMBED
   (4)   RECORD FORMAT            (16)  REPLICATE
   (5)   LOGICAL RECORD LENGTH    (17)  CONTROL INTERVAL SIZE DATA
   (6)   KEY LENGTH               (18)  % FREESPACE CONTROL AREA
   (7)   KEY OFFSET               (19)  % FREESPACE CONTROL INTERVAL
   (8)   AVGREC                   (20)  SHAREOPTIONS CROSSREGION
   (9)   AVG VALUE                (21)  SHAREOPTIONS CROSSSYSTEM
  (10)   SPACE PRIMARY            (22)  LAST MODIFIED USERID
  (11)   SPACE SECONDARY          (23)  LAST DATE MODIFIED
  (12)   SPACE DIRECTORY          (24)  LAST TIME MODIFIED

USE ENTER TO PERFORM SORT;
USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 60. Specifying Data Class Sort Criteria

You need to specify either an ascending or descending order sequence for each of the sorting fields. ISMF saves the sorting criteria for the life of the application, or until you update the criteria. If you specify N, no, in the RESPECIFY SORT CRITERIA field, ISMF uses the most recent sorting order.

Based on the sorting criteria, ISMF displays a list of data classes like the one shown in Figure 61. On the listing, you can enter line operators such as COPY, ALTER, DELETE, DISPLAY, and HIDE against individual entries.

```
                          DATA CLASS LIST
COMMAND ===>                                       SCROLL ===> CSR
                                               Entries 1-1 of 1
                                               Data Columns 3-9 of 24


CDS NAME:  SMS.SCDS1.SCDS

ENTER LINE OPERATORS BELOW:


   LINE       DATACLAS                                            AVG
   OPERATOR   NAME      RECORG  RECFM  LRECL  KEYLEN  KEYOFF  AVGREC  VALUE
   ---(1)---- --(2)---  -(3)--  -(4)-  -(5)-  -(6)--  -(7)--  -(8)--  -(9)-
              DATAV     --      VB     255    ---     -----   U       255
   ---------- --------  ------  BOTTOM OF  DATA  ------  --------  ----------











USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.
```

Figure 61. Displaying a List of Data Classes

You can scroll up, down, left, and right. However, ISMF always displays the LINE OPERATOR and DATACLAS NAME fields on the left side of the panel.

Altering SMS Classes and Storage Groups

When you want to change an SMS class or storage group definition, you can enter the ALTER line operator against an entry in a DATA SET LIST (option 1 in the ISMF Primary Option Menu) or you can select the ALTER option from one of the SMS class or storage group applications. Note that you cannot alter anything in the active configuration. After altering SMS classes and storage groups, you must activate the SCDS for the changes to affect storage management.

You can modify storage groups, management classes, storage classes, and data classes using the ALTER option.

- You can use the ALTER line operator to change which management or storage class is associated with a data set.

- You can use the ISMF ALTER dialogs to change the definition of an SMS class or storage group.

Altering Storage Groups

As the needs of your installation change, you might find the need to modify storage groups. To modify a storage group, go to the Storage Group Application Selection panel, shown in Figure 8 on page 26, and specify the name of the SCDS containing the storage group you want to modify. Specify the storage group name and select option 3, ALTER. ISMF displays the Storage Group Alter panel.

Redefining Volumes among Storage Groups

If you have defined a volume to belong to one storage group but now want it to belong to a different storage group, go to the Storage Group Application Selection panel. Specify the CDS name, and storage group name. Press ENTER to get the Storage Group Volume Selection panel shown in Figure 62 on page 97.

```
┌─────────────────────────────────────────────────────────────────────────┐
│                      STORAGE GROUP VOLUME SELECTION                       │
│ COMMAND ===>                                                              │
│ CDS NAME:           SMS.SCDS1.SCDS                                        │
│                                                                           │
│ STORAGE GROUP NAME: PRIMARY                                               │
│ STORAGE GROUP TYPE: POOL                                                  │
│                                                                           │
│ SELECT ONE OF THE FOLLOWING OPTIONS   ===> 4                             │
│                                                                           │
│    1  DISPLAY      - Display SMS Volume Status (Pool only)               │
│    2  DEFINE       - Add a Volume to Volume Serial Number List           │
│    3  ALTER        - Alter Specific Volume and Statuses   (Pool only)    │
│    4  DELETE       - Delete a Volume from Volume Serial Number List      │
│                                                                           │
│ SPECIFY VOLUME SERIAL NUMBER            ===> 123456                       │
│                                                                           │
│                                                                           │
│                                                                           │
│                                                                           │
│                                                                           │
│ USE ENTER TO PERFORM SELECTION;                                          │
│ USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.                      │
└─────────────────────────────────────────────────────────────────────────┘
```

Figure 62. Moving a Volume from One Storage Group to Another

In the SPECIFY VOLUME SERIAL NUMBER field, enter the serial number of the
volume you want to redefine, select option 4 to delete it from its current storage
group, and press ENTER. Then, return to the Storage Group Application
Selection panel, enter the name of the new storage group to contain the
volume, and select option 4, VOLUME, which returns you to the Storage Group
Volume Selection panel. Specify the volume serial number, select option 2 to
define it to the storage group, and press ENTER.

You are responsible for maintaining consistent physical connections and being
aware of multivolume data sets. You must also maintain storage group and
management class consistency. If a data set has a management class with
certain migration and backup attributes, and you redefine the volume containing
that data set from one storage group to another, the new storage group might
not be eligible for the same migration and backup processing. You must also
maintain storage group and storage class consistency.

Altering Management Classes

As the needs of your installation change, you might find the need to modify
management classes. To modify a management class, go to the Management
Class Application Selection panel, shown in Figure 15 on page 38, and specify
the name of the SCDS containing the management class you want to change.
Specify the management class name and select option 4, ALTER. ISMF dis-
plays the Management Class Alter panel shown in Figure 63 on page 98.

```
┌─────────────────────────────────────────────────────────────────────────┐
│                       MANAGEMENT CLASS ALTER              Page 1 of 3     │
│  COMMAND ===>                                                             │
│                                                                          │
│  SCDS NAME:             SMS.SCDS1.SCDS                                    │
│  MANAGEMENT CLASS NAME:  NOACTION                                        │
│                                                                          │
│  TO ALTER MANAGEMENT CLASS, SPECIFY:                                     │
│                                                                          │
│    DESCRIPTION ===>                                                      │
│                ===>                                                      │
│                                                                          │
│    EXPIRATION ATTRIBUTES                                                  │
│      EXPIRE AFTER DAYS NON-USAGE ===> NOLIMIT     (1 to 9999 or NOLIMIT) │
│      EXPIRE AFTER DATE/DAYS      ===> NOLIMIT     (0 to 9999, yyyy/mm/dd │
│                                                    or NOLIMIT)           │
│                                                                          │
│    RETENTION LIMIT               ===> NOLIMIT     (0 to 9999 or NOLIMIT) │
│                                                                          │
│                                                                          │
│                                                                          │
│  USE ENTER TO PERFORM VERIFICATION; USE DOWN COMMAND TO VIEW NEXT PANEL; │
│  USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT. │
└─────────────────────────────────────────────────────────────────────────┘
```

Figure 63. Altering a Management Class

This panel contains the same parameters as the Management Class Define
panel. To update the management class, you can modify parameters on this
and the two following ALTER panels. You can then enter the END command to
save your changes and return to the Management Class Application Selection
panel. Existing data sets are affected at the next DFHSM automatic function or
DFHSM volume or data set command function. RETENTION LIMIT processing is
not performed retroactively on existing data sets.

Altering Storage Classes

As the needs of your complex change and as new hardware becomes available,
you might experience a need to change the storage classes. To modify a
storage class, go to the Storage Class Application Selection panel, shown in
Figure 20 on page 48, and specify the name of the SCDS containing the storage
class you want to change. Specify the name of the storage class and select
option 4, ALTER. ISMF displays the Storage Class Alter panel shown in
Figure 64 on page 99.

```
                           STORAGE CLASS ALTER
COMMAND ===>

SCDS NAME:              SMS.SCDS1.SCDS
STORAGE CLASS NAME: DBFAST

TO ALTER STORAGE CLASS, SPECIFY:

  DESCRIPTION ===> THIS DATA BASE STORAGE CLASS PROVIDES THE FASTEST
              ===> RESPONSE TIME OF ALL THE DATA BASE STORAGE CLASSES.

  PERFORMANCE OBJECTIVES

  DIRECT MILLISECOND RESPONSE      ===> 5        (1 to 999 or blank)
  DIRECT BIAS                      ===> R        (R, W or blank)

  SEQUENTIAL MILLISECOND RESPONSE ===> 5         (1 to 999 or blank)
  SEQUENTIAL BIAS                  ===> R        (R, W or blank)

  AVAILABILITY                     ===> STANDARD  (STANDARD or CONTINUOUS)
  GUARANTEED SPACE                 ===> Y         (Y or N)

USE ENTER TO PERFORM VERIFICATION;
USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT.
```

Figure 64. Altering a Storage Class

This panel contains the same parameters as the Storage Class Define panel.
To update the storage class, modify the parameters on this panel. You can
then enter the END command to save the changes and return to the Storage
Class Application Selection panel. After activating the SCDS, the changes
apply to new data sets for volume selection and to existing data sets when
determining if a data set to be accessed will be cached.

Altering Data Classes

As the needs of your installation change, you might find the need to change
data classes. To modify a data class, go to the Data Class Application
Selection panel, shown in Figure 22 on page 54, and specify the name of the
SCDS containing the data class you want to change. Specify the data class
name and select option 4, ALTER. ISMF displays the Data Class Alter panel
shown in Figure 65 on page 100.

```
┌─────────────────────────────────────────────────────────────────────────┐
│                          DATA CLASS ALTER                    Page 1 of 2  │
│   COMMAND ===>                                                            │
│                                                                          │
│   SCDS NAME:       SMS.SCDS1.SCDS                                         │
│   DATA CLASS NAME: DATAV                                                  │
│                                                                          │
│   TO ALTER DATA CLASS, SPECIFY:                                          │
│                                                                          │
│     DESCRIPTION ===> VARIABLE LENGTH RECORD FORMAT FOR SCRIPT DATA SETS,  │
│                 ===> CLISTS, AND PROGRAM SOURCE CODE.                     │
│     RECORG         ===>                 (KS, ES, RR, LS or blank)         │
│     RECFM          ===> VB              (any valid RECFM combination or blank) │
│     LRECL          ===> 255             (1 to 32761 or blank)             │
│     KEYLEN         ===>                 (0 to 255 or blank)               │
│     KEYOFF         ===>                 (0 to 32760 or blank)             │
│     SPACE AVGREC   ===> U               (U, K, M or blank)                │
│           AVG VALUE ===> 255            (0 to 65535 or blank)             │
│           PRIMARY   ===> 5000           (0 to 999999 or blank)            │
│           SECONDARY ===> 5000           (0 to 999999 or blank)            │
│           DIRECTORY ===>                (0 to 999999 or blank)            │
│     RETPD OR EXPDT  ===>                (0 to 9999, yyyy/mm/dd or blank)   │
│                                                                          │
│   USE ENTER TO PERFORM VERIFICATION; USE DOWN COMMAND TO VIEW NEXT PANEL; │
│   USE HELP COMMAND FOR HELP; USE END COMMAND TO SAVE AND EXIT; CANCEL TO EXIT. │
└─────────────────────────────────────────────────────────────────────────┘
```

Figure 65. Altering a Data Class

This panel contains the same parameters as the Data Class Define panel. To update the data class, you can modify parameters on both this and the second page of the Data Class Alter panel. You can then enter the END command to save the changes and return to the Data Class Application Selection panel. After activating the SCDS, changes apply only to new data sets. Existing data sets are not affected.

Copying SMS Classes and Storage Groups

To simplify the creation of new SMS classes, storage groups, and SCDS, you can copy existing ones and modify them. In the LINE OPERATOR field of any List panel, you can enter the COPY line operator and press ENTER to duplicate an existing SMS class or storage group. However, for storage groups, the volume list associated with the storage group is not copied. You receive the Copy Entry Panel shown in Figure 66 on page 101.

```
┌─────────────────────────────────────────────────────────────────────────┐
│                         COPY ENTRY PANEL                                  │
│  COMMAND ===>                                                             │
│                                                                           │
│   CONSTRUCT WILL BE COPIED FROM:                                          │
│                                                                           │
│    DATA SET NAME:   'SMS.SCDS1.SCDS'                                       │
│    CONSTRUCT NAME: DATAV                                                   │
│    CONSTRUCT TYPE: DATA CLASS                                             │
│                                                                           │
│   SPECIFY "COPY TO" CONSTRUCT:                                            │
│                                                                           │
│     DATA SET NAME ===> 'SMS.SCDS.SCDS2'                                    │
│                                  (1 to 44 characters)                      │
│                                                                           │
│     CONSTRUCT NAME ===>          (1 to 8 characters, fully specified)      │
│                                                                           │
│   REPLACE LIKE-NAMED CONSTRUCT ===> N  (Y or N)                            │
│                                                                           │
│                                                                           │
│     PERFORM ALTER            ===> N  (Y or N)                              │
│                                                                           │
│  USE ENTER TO PERFORM COPY;                                               │
│  USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.                       │
└─────────────────────────────────────────────────────────────────────────┘
```

Figure 66. Copying SMS Classes and Storage Groups

The *from* DATA SET NAME field identifies the source that you are copying. It is
primed with the value you specified on the Application Selection panel. The
from CONSTRUCT NAME field identifies the name of the SMS class or storage
group you are copying. It is primed with the value from the NAME field of the
List panel. The *from* CONSTRUCT TYPE field identifies whether you are copying
a data class, storage class, management class, or storage group. It is primed
with the type of List panel from which you issued the COPY command.

The *to* DATA SET NAME field identifies the target of the copy. It must be the
name of an SCDS. It is primed with the value of the *from* DATA SET NAME if
the *from* DATA SET NAME contains an SCDS name. It is primed with blanks if
the *from* DATA SET NAME is 'ACTIVE'. The *to* CONSTRUCT NAME field identi-
fies the name of the SMS class or storage group copy. It is primed with blanks.
The copy has the same *type* as the original.

If the SMS class or storage group that you identify in the *to* fields already
exists, you specify whether or not you want to replace it in the REPLACE
LIKE-NAMED CONSTRUCT field. If you indicate Y, yes, the replacement occurs.
If you indicate N, no, the copy does not occur. The field is primed with N, which
is the default.

In the PERFORM ALTER field, you indicate if you want to change some of the
attributes of the copy you are creating. If you specify Y, yes, and the copy exe-
cutes successfully, you receive the pertinent Alter panel. If you specify N, no,
you remain on the Copy Entry Panel, where you can perform another copy or
return to the original List panel.

When you have specified the values, press ENTER to perform the copy.

Deleting Volumes from Storage Groups

When deleting a volume from a storage group, first set the volume status to DISNEW. This prevents any new allocations to the volume, while allowing existing data sets to still be accessed with DISP=SHR or OLD. Next, if the volume is to be reused as a non-SMS volume, run a DFDSS NONSMS CONVERTV to convert it out of DFSMS. Finally, remove the volume from the storage group. Be careful when moving or removing a volume from a storage group because the volume could contain part of a multivolume data set.

To remove a volume from a storage group, select the Storage Group Application Selection panel. Specify the CDS name and the storage group name, that contains the volume you want to delete. Press ENTER and you will get the Storage Group Volume Selection panel shown in Figure 62 on page 97. From the Storage Group Volume Selection panel, indicate the individual volume you want to delete in the SPECIFY VOLUME SERIAL NUMBER field, select option 4, and press ENTER. The changes take effect when you activate this updated configuration.

After deletion, the volume is only eligible for non-SMS allocations. However, if you are reassigning the volume to another storage group, you need to define the volume to the storage group and then activate the updated configuration.

Deleting Storage Groups

You might want to delete a pool type storage group so that you can define all of its volumes to other storage groups or so that you can use the volumes for non-SMS allocations.

If you plan to use the volumes for non-SMS allocations, you need to delete or move all SMS-managed data sets. The following procedure shows how to do this and how to remove the storage group from the SMS configuration:

1. Set the storage group status to DISNEW to prevent any further new allocations to volumes defined to the storage group. As time passes, DFHSM space management processing clears most of the data sets off the volumes in the storage group you are deleting, provided AUTO MIGRATE is 'Y' and most of the data sets have an associated management class.

2. Many data sets might linger on the volumes. When the volumes are sufficiently empty, move the remaining data sets off of the storage group volumes using DFDSS.

3. Build an SCDS to contain the new configuration. Remove the storage group definition from the SCDS and change the storage group ACS routine so that it does not select the deleted storage group.

4. Activate the new configuration.

To make volumes eligible for non-SMS allocations, you need to INIT them through ICKDSF to non-SMS or CONVERT them through DFDSS to non-SMS.

When you delete a VIO or dummy type storage group, you do not have to worry about redefining volumes to other storage groups, but you might need to modify the storage group ACS routine.

Deleting SMS Classes

You need to exercise caution when deleting storage classes and management classes, because a user or system might attempt to use a data set that references them.

For example, assume a data set has some specified migration attributes, and at some point it meets the migration criteria. After the data set has been migrated, you delete its associated management class. If the data set is recalled in the future, it must go through the management class ACS routine. If the ACS routine fails to override the undefined or deleted management class, the recall fails because a management class that does not exist is associated with the data set. If this occurs, you must either rewrite the management class ACS routine or specify that the data set should be recalled as a non-SMS data set. If you rewrite the ACS routine, you must check for the deleted management class and assign a valid management class to the data set. If you decide to bypass ACS processing, be aware that the data set is recalled to non-SMS storage.

Another consequence of deleting management classes involves DFHSM automatic processing. When DFHSM runs through its automatic cycle and processes data sets based on their management class attributes, it attempts to retrieve nonexistent management class definitions. Consequently, DFHSM skips the processing of these data sets.

Instead of deleting storage classes and management classes, you should prevent any new allocation from using them by rewriting the corresponding ACS routine to override the deleted storage class or management class. It is safer and provides the possibility that all data sets that reference the management class or storage class might eventually be overridden. Also, you can delete and use DFDSS to identify all data sets with storage class or management class then use ALTER to change to a new name. If you decide to delete a storage class or management class, make certain that you inform all users well in advance.

Because data classes are used only at data set creation, they do not have these problems. The original data class name is never referenced or reused.

If you decide to delete an SMS class, enter the following information on the pertinent Class Application Selection panel:

```
SCDS NAME          ===> SMS.SCDS1.SCDS
xxxx CLASS NAME    ===> *
SELECT OPTION      ===> 1
```

After you press ENTER, ISMF will display the Class List panel. Next to the SMS class that you want to delete, enter DELETE in the LINE OPERATOR field. When you press ENTER, you will see the Confirm Delete Request panel shown in Figure 67 on page 104.

```
┌─────────────────────────────────────────────────────────────────┐
│                      CONFIRM DELETE REQUEST                       │
│ COMMAND ===>                                                      │
│                                                                   │
│    TO CONFIRM DELETION ON THE FOLLOWING ELEMENT:                  │
│                                                                   │
│         ELEMENT NAME:      DATAV                                   │
│         ELEMENT TYPE:      DATA CLASS                              │
│         RESIDING IN SCDS:  SMS.SCDS1.SCDS                          │
│                                                                   │
│                                                                   │
│    SPECIFY THE FOLLOWING:                                         │
│         PERFORM DELETION ===> N          (Y or N)                 │
│                                                                   │
│                                                                   │
│                                                                   │
│                                                                   │
│                                                                   │
│                                                                   │
│                                                                   │
│                                                                   │
│    USE ENTER TO PERFORM OPERATION;                                │
│    USE HELP COMMAND FOR HELP; USE END COMMAND TO EXIT.            │
└─────────────────────────────────────────────────────────────────┘
```

Figure 67. Deleting an SMS Class

Confirm that the displayed SMS class is the one that you want to delete. If it is, enter Y for yes and press ENTER. The SMS Class List should appear with '*DELETE' in the LINE OPERATOR field next to the deleted SMS class.

After deleting SMS classes, activate the modified SMS configuration to make your changes part of the active configuration.

Converting Volumes to SMS

You can convert empty volumes through ISMF using the INIT line operator from the Volume Application. You can convert previously used non-SMS volumes using the CONVERTV command of DFDSS. You can also unconvert SMS volumes using the CONVERTV command of DFDSS.

See *MVS SML: Storage Management Subsystem Migration Planning Guide* for an explanation of how to convert and unconvert volumes.

Chapter 11. Recovering SMS Information

Recovering Control Data Sets

You should allocate a spare ACDS and COMMDS when you allocate your originals. Having spare copies facilitates the recovery process if SMS cannot access the originals because of I/O errors. Place the spares on a device that every system in your complex can access. Make certain that they reside on a different device than your originals.

Recovering the SCDS

For purposes of recovery, treat an SCDS the same as any other source VSAM linear data set. You can use DFHSM to manage its availability, which relieves you from having to allocate any spares.

Recovering the ACDS

You can recover from errors that prevent access to the ACDS if you have allocated a spare. All permanent errors that make the ACDS unreadable or unwriteable require intervention. For permanent I/O errors to the ACDS, the messages IGD041I and IGD040D appear on the operator console. You have the option of re-trying the failing operation or using the spare ACDS.

If you decide to use the spare ACDS, you need to perform three steps. First, reply 'C' on the operator console. Next, copy the data from the SMS address space to the spare ACDS by issuing the following command from a system in the SMS complex:

SETSMS SAVEACDS(spare.acds)

Finally, tell the system to use the spare ACDS by issuing the following command:

SETSMS ACDS(spare.acds)

You only need to issue the SETSMS ACDS(spare.acds) command on one system. The COMMDS is updated to reflect this change. As the other systems in the SMS complex access the COMMDS (based on the INTERVAL value in their respective IGDSMSmm members), they automatically switch to the new ACDS.

Recovering the COMMDS

If you cannot access the COMMDS, you can recover from the error if you allocated a spare data set. All permanent errors that make the COMMDS unreadable require intervention. For permanent I/O errors to the COMMDS, the messages IGD041I and IGD070D appear on the operator console. Reply 'S' on the operator console and issue the following command from a system in the SMS complex:

SETSMS COMMDS(spare.commds)

One of three situations will result.

If the spare COMMDS is empty it gets formatted automatically, and SMS writes the in-storage copy of the current COMMDS into the spare.commds. You then

need to issue the following command on all the remaining systems in the SMS complex:

```
SETSMS COMMDS(spare.commds)
```

If the spare COMMDS is not empty but describes an ACDS that is not currently active in the SMS complex, then SMS issues the message IGD076D. This message asks if you want to use the contents of the COMMDS and the ACDS to which it points. Reply 'C' to cause SMS to replace the contents of the spare.commds with the in-storage copy of the current COMMDS. You then need to issue the following command on all the remaining systems in the SMS complex:

```
SETSMS COMMDS(spare.commds)
```

If the spare COMMDS is not empty but describes the ACDS that is currently active in the SMS complex, you need to issue the following command on all the remaining systems in the SMS complex:

```
SETSMS COMMDS(spare.commds)
```

Without a usable COMMDS, the systems in the SMS complex have no means of communication. Other systems in the SMS complex are aware of the error, but they are unaware of the switch to a new COMMDS until you inform them.

If you can access the current COMMDS but you want to use an alternate one, you only need to issue the SETSMS command from one system. The other systems in the SMS complex detect the change from the old COMMDS to the new COMMDS and they automatically switch to the new one.

Notes on Recovering ACDS and COMMDS

If an SMS system goes down (a temporary condition) or has not yet been started with SMS, then it is not part of the system-managed storage environment. When you activate SMS on the system, it uses the ACDS and COMMDS that are specified in its IGDSMSmm member. If the COMMDS and ACDS are different than the ones used by the remainder of the SMS complex, then the system runs as a separate SMS complex.

Consequently, when you change the current ACDS or COMMDS, update the IGDSMSmm member of SYS1.PARMLIB for each system in the SMS complex.

Recovering the SMS Address Space

If the SMS address space fails, SMS automatically attempts to restart up to six times. If the system fails to recover the address space after the sixth restart, you have two options:

1. You can terminate the SMS address space and then restart SMS using the T SMS=xx command. If you have cross memory interactions, users' address spaces wait for the SMS address space to restart.

2. You can allow SMS to attempt another restart by replying to the system-generated message:

   IGD032D
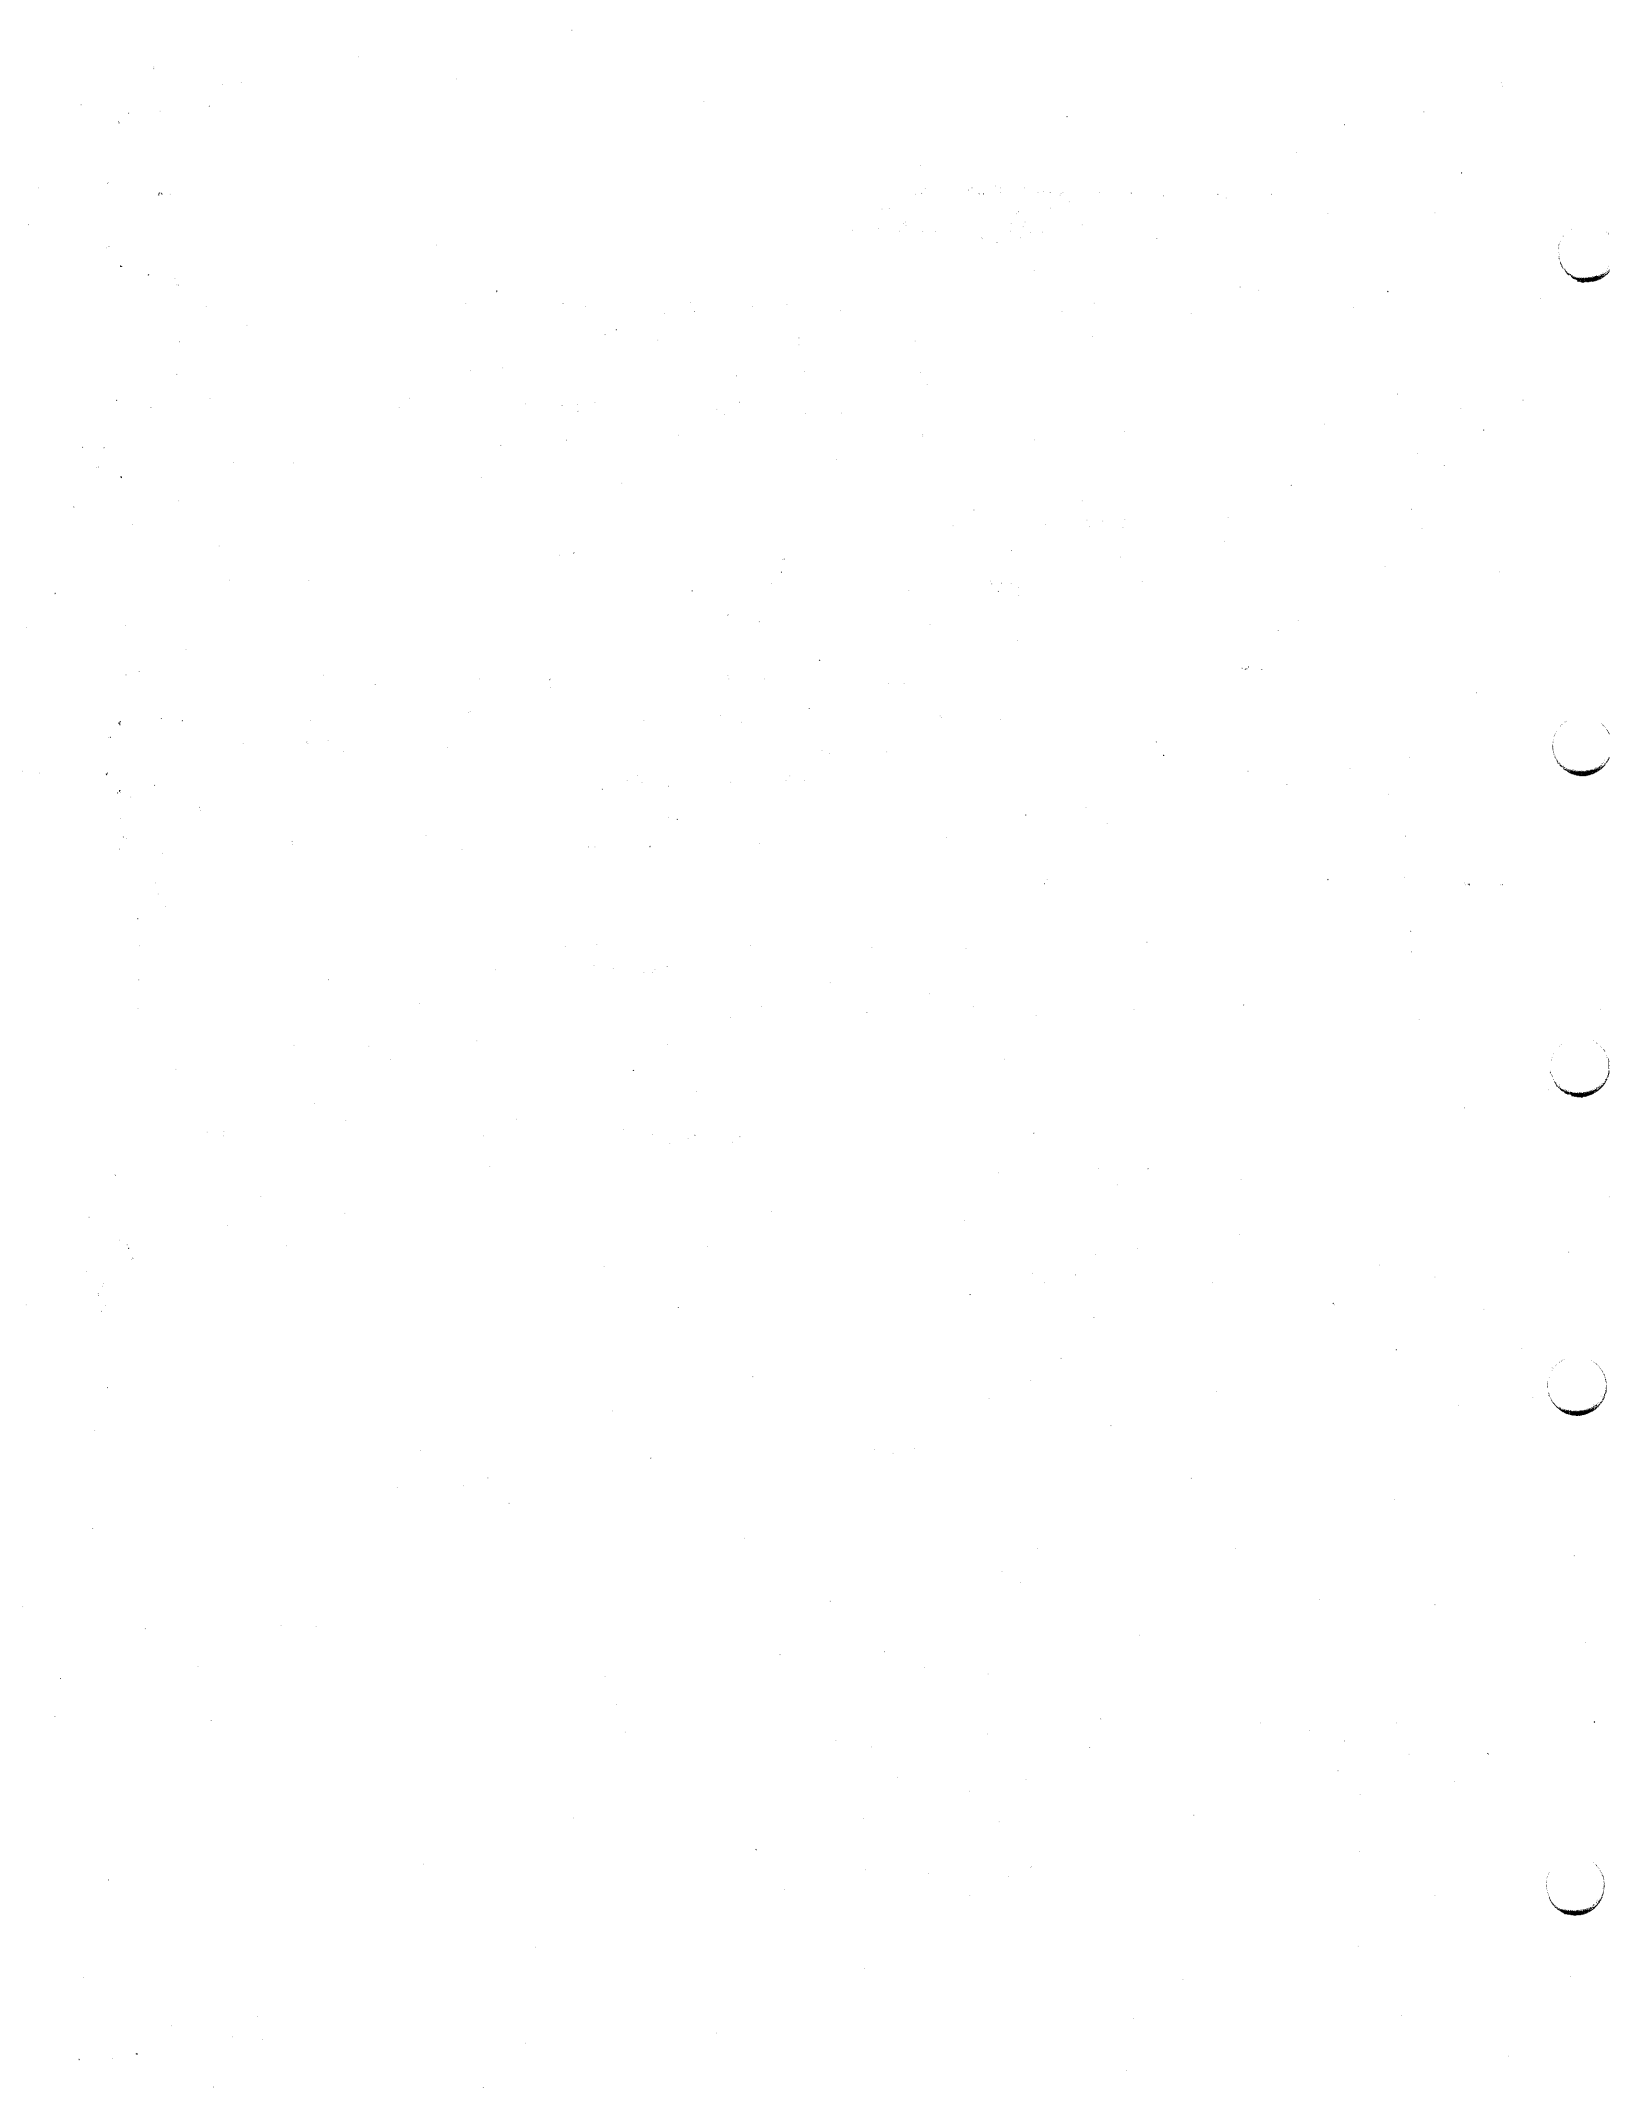
Chapter 12. Protecting SMS

To ensure that SMS operates correctly, you need to prevent unauthorized end users from modifying information in certain control data sets, such as SMS class and storage group definitions. Some ISMF commands and functions also require protection. This chapter explains how to use RACF to establish authorization levels for protecting these data sets, commands, and functions.

RACF and SMS Interactions

The **resource owner** is the actual owner of the data set covered by the RACF DATASET profile. The resource owner is extracted based on data set name. If the resource owner is not identified, the high-level qualifier is returned.

The extracted resource owner can be used to extract the default SMS classes and application identifier. If the resource owner is a user and no default SMS information is available from the user profile, the default information from the group profile is used. If the resource owner is a group, then the defaults for the group profile is used. You can protect the ability to update the resource owner field, RESOWNER, in the data set RACF profile. See also "FIELD Resource Class" on page 115.

After ACS routines have been run, RACF is invoked to verify the user's authority to create the data set (CREATE/ALTER) and the resource owner's authority to use the STORCLAS and MGMTCLAS (READ). You can protect the ability of a resource owner to use management class and storage class through STORCLAS/MGMTCLAS Resource class.

See *MVS SML: Managing Data Sets* for additional information on using RACF in an SMS environment.

Protecting ISMF Functions

You can use RACF authorization to protect the following categories of ISMF functions:

1. The entire ISMF component

2. The individual ISMF applications:

> Profile
> Data Set
> Volume
> Management Class
> Data Class
> Storage Class
> Storage Group
> Automatic Class Selection
> Control Data Set

3. The ISMF line operators

4. The ISMF commands

ISMF relies on the **RACF program control feature** to protect many of its applications. The program control feature prevents unauthorized end users from running selected ISMF programs. To use the feature, you must activate the RACF Program Class and define your selected ISMF programs to RACF.

With program control you can set up authorization levels for each of these categories, varying the level within a particular category to suit the needs of your installation.

See *Resource Access Control Facility (RACF) Security Administrator's Guide* for a detailed description of how to use the RACF program control features.

Establishing an Authorization Structure

The RACF program control performs authorization checking before it allows access to an ISMF function. The authorization level of a function's load module determines its degree of protection. Individual end users can execute an ISMF function if one of the following conditions is true:

- They are authorized to execute the corresponding load module.
- Their RACF profile contains the OPERATIONS attribute.
- Their group is authorized to execute the load module.
- RACF is disabled or the control program is turned off.
- The universal access authority (UACC) for the load module is READ or greater, making the load module available to anyone who can access ISMF.

Locating Module Names for ISMF Applications

Figure 68 lists the load module names you can use to limit access to certain ISMF applications. The names are included in the panel coding for the ISMF Primary Option Menu for Storage Administrators, which is stored in the panel library with a member name of DGTSMMD2. If you want to limit access to all of ISMF, use module name DGTFMD01.

Figure 68. Module Names for ISMF Applications

Application	Module Name
ISMF profile	DGTFPF00
Data set	DGTFDS00
Volume	DGTFVA00
Management class	DGTFMCCD
Data class	DGTFDCCD
Storage class	DGTFSCCD
Storage group	DGTFSGDR
Automatic class selection	DGTFFLAD
Control data set	DGTFSACD

Locating Module Names for ISMF Functions

Restricting access to the DGTFPF05 module prevents end users from gaining access to the Primary Option Menu for Storage Administrators. Figure 69 lists this and the other load module names you can use to limit access to certain ISMF functions.

Figure 69. Module Names for ISMF Functions

Function	Module Name
User mode	DGTFPF05
Logging and abend control	DGTFPF02
ISMF job statement information	DGTFPF03
DFDSS execute statement information	DGTFPF04
ICKDSF execute statement information	DGTFPF20
Data set print execute statement information	DGTFPF21
IDCAMS execute statement information	DGTFPF22
Data class DEFINE	DGTFDCDA
Data class ALTER	DGTFDCAA
Data class DISPLAY	DGTFDCDI
Data class LIST	DGTFDCLD
Storage class DEFINE	DGTFSCDA
Storage class ALTER	DGTFSCAA
Storage class DISPLAY	DGTFSCDI
Storage class LIST	DGTFSCLD
Management class DEFINE	DGTFMCDA
Management class ALTER	DGTFMCAA
Management class DISPLAY	DGTFMCDI
Management class LIST	DGTFMCLD

Locating Module Names for Line Operators and Commands

Figure 70 and Figure 71 on page 112 list the module names for ISMF line operators and commands. The names are stored in command tables in the DGTLLIB load library.

Figure 70 (Page 1 of 2). Module Names for ISMF Line Operators

Line Operator	Data Set	Volume	Data Class	Management Class	Storage Class	Storage Group
ALTER	DGTFAL01	-	DGTFALD1	DGTFALM1	DGTFALS1	DGTFALG1
ANALYZE	-	DGTFAZ01	-	-	-	DGTFAZ01
BROWSE	DGTFBR01	-	-	-	-	-
CLIST	DGTFCL01	DGTFCL01	-	-	-	-
COMPRESS	DGTFCM01	DGTFCS01	-	-	-	-
CONDENSE	DFQFCND1	-	-	-	-	-
CONVERTV	-	DGTFCN01	-	-	-	-
COPY	DGTFCY01	DGTFCV01	DGTFCAD1	DGTFCAM1	DGTFCAS1	DGTFCAG1
DEFRAG	-	DGTFDF01	-	-	-	-
DELETE	DGTFDL01	-	DGTFDND1	DGTFDNM1	DGTFDNS1	DGTFDNG1
DISPLAY	-	-	DGTFDID1	DGTFDIM1	DGTFDIS1	-
DUMP	DGTFDP01	DGTFDM01	-	-	-	-
EDIT	DGTFED01	-	-	-	-	-
HALTERDS	DFQFHA01	-	-	-	-	-
HBACKDS	DFQFHB01	-	-	-	-	-

Figure 70 (Page 2 of 2). Module Names for ISMF Line Operators

Line Operator	Data Set	Volume	Data Class	Management Class	Storage Class	Storage Group
HBDELETE	DFQFHBD1	-	-	-	-	-
HDELETE	DFQFHD01	-	-	-	-	-
HIDE	DGTFHI01	DGTFHI01	DGTFHI01	DGTFHI01	DGTFHI01	DGTFHI01
HMIGRATE	DFQFHM01	-	-	-	-	-
HRECALL	DFQFHRL1	-	-	-	-	-
HRECOVER	DFQFHRC1	-	-	-	-	-
INIT	-	DGTFIN01	-	-	-	DGTFIN01
INSPECT	-	DGTFIV01	-	-	-	DGTFIV01
LISTSYS	-	-	-	-	-	DGTFLIC1
LISTVOL	-	-	-	-	-	DGTFLVC1
MESSAGE	DGTFMS00	DGTFMS00	DGTFMS00	DGTFMS00	DGTFMS00	DGTFMS00
RAUTH	-	DGTFRA01	-	-	-	-
RELEASE	DGTFRL01	DGTFRV01	-	-	-	-
RESTORE	DGTFRT01	DGTFRO01	-	-	-	-
REFORMAT	-	DGTFRF01	-	-	-	DGTFRF01
SECURITY	DGTFSRD1	-	-	DGTFSRD1	DGTFSRD1	-
SETCACHE	-	DGTFCB01	-	-	-	-
SORTREC	ICEPSR01	-	-	-	-	-

Figure 71. Module Names for ISMF Commands

Command	Module Name
ACTIVATE	DGTFACAT
CANCEL	-
CLEAR	DGTFCR01
COMPRESS	DGTFCP01
COPY	DGTFCO01
DOWN	DGTFDO01
DUMP	DGTFDU01
END	-
ERTB	DGTFER02
FILTER	DGTFFI01
FIND	DGTFFN01
FOLD	DGTFFU01
HELP	-
LEFT	DGTFLE01
PROFILE	DGTFPF01
REFRESH	-
RELEASE	DGTFRE01
RESHOW	DGTFRW01
RESTORE	DGTFRR00
RETURN	-
RIGHT	DGTFRI01
SAVE	DGTFSLDS
SORT	DGTFSO01
UP	DGTFUP01
VALIDATE	DGTFVLVA

If you want to look at the command tables, you need the data set name that
your installation uses for the load library. The install process for ISMF puts the
load modules in SYS1.DGTLLIB (for MVS/DFP/ISMF and DFDSS/ISMF), and
SYS1.DFQLLIB (for DFHSM/ISMF). However, your installation's post-install pro-
cedures might involve moving the ISMF modules to other libraries. If the

libraries have moved but have kept the same names, you can determine their data sets names by issuing the TSO/E LISTALC command and scanning the low-level qualifiers for DGTLLIB and DFQLLIB.

Protecting Modules

Take the following three program control steps to protect modules:

1. Use the RDEFINE command or the RACF ISPF entry panels to identify the modules you want to protect. To define the modules to RACF, supply the name of the load module that you want to protect, the name of the data set that contains the load module, and the volume serial number of the volume that contains the data set. RACF will add each module that you identify to the profile for the PROGRAM general resource class.

   When you define the modules, you have several options:

   - If you want to define several modules at the same time, you can use asterisk notation. For example, DGT* represents all of the modules beginning with the letters DGT.

   - You can add an access list with userids and/or group names and their associated access authority to the profile.

   - You can define the UACC to give default access to all users or to none.

   - You can use the AUDIT parameter to set up or to bypass RACF logging.

2. Use the PERMIT command to allow end users to execute an application, line operator, or command associated with a module.

3. To prevent unauthorized users from copying a program, renaming it to a name that is unknown to the program control, and then executing the renamed program, you should protect the PDS libraries containing RACF-controlled programs with a UACC of NONE. In order for users to execute programs in these libraries, place the libraries in the LNKLIST con- catenation. See *RACF Security Administrator's Guide* for more information.

FACILITY Resource Class

To prevent unauthorized users from copying a program, renaming it to a name that is unknown to the program control, and then executing the renamed program, you should protect the PDS libraries containing RACF controlled pro- grams with a UACC of NONE. In order for users to execute programs in these libraries, place the libraries in the LNKLIST concatenation. See *RACF Security Administrator's Guide* for additional details.

You can use the RACF FACILITY resource class to control the execution of func- tions, keywords, and commands beyond the load module level. It can control the ability to:

- Activate a configuration
- Perform catalog functions on SMS data sets
- Perform DFDSS functions on SMS data set.

Catalog and DFDSS functions on SMS data sets are allowed if profiles do not exist or if FACILITY resource class is not active. Also, the user must be APF authorized.

You need to define the following RACF profiles:

STGADMIN.IGD.ACTIVATE.CONFIGURATION
> Protects the ability to activate a configuration.

STGADMIN.IDC.DIAGNOSE.CATALOG
> Protects the ability to run the access method services DIAGNOSE command against catalogs.

STGADMIN.IDC.DIAGNOSE.VVDS
> Protects the ability to run the access method services DIAGNOSE command against a VVDS when a comparison against the BCS is performed. In this case, the BCS is protected.

STGADMIN.IGG.ALTBCS
> Protects the ability to alter the BCS entry only.

STGADMIN.IGG.ALTER.UNCONVERT
> Protects the ability to alter a VSAM data set from SMS to non-SMS.

STGADMIN.IGG.ALTER.SMS
> Protects the ability to alter the storage class and management class of the data sets.

STGADMIN.IGG.DIRCAT
> Protects the ability to use directed catalog.

STGADMIN.IGG.DEFNVSAM.NOBCS
> Protects the ability to define a non-VSAM entry without a BCS entry.

STGADMIN.IGG.DEFNVSAM.NONVR
> Protects the ability to define a non-VSAM entry without an NVR entry.

STGADMIN.IGG.DELNVR.NOBCSCHK
> Protects the ability to delete an NVR record without checking a BCS entry.

STGADMIN.IGG.DELETE.NOSCRTCH
> Protects the ability to delete a BCS entry without deleting the VTOC entry and NVR.

STGADMIN.IGG.DELGDG.FORCE
> Protects the ability to delete a GDG with the FORCE operand.

STGADMIN.ADR.COPY.BYPASSACS
> Protects the ability to have DFDSS COPY bypass the ACS routines and use the original or user-specified storage class and management class.

STGADMIN.ADR.RESTORE.BYPASSACS
> Protects the ability to have DFDSS RESTORE bypass the ACS routines and use the original or user-specified storage class and management class.

STGADMIN.ADR.COPY.INCAT
> Protects the ability to use the INCAT parameter on a DFDSS COPY.

STGADMIN.ADR.DUMP.INCAT
> Protects the ability to use the INCAT parameter on a DFDSS DUMP.

STGADMIN.ADR.RESTORE.INCAT
> Protects the ability to use the INCAT parameter on a DFDSS RESTORE.

STGADMIN.ADR.CONVERTV.INCAT
> Protects the ability to use the INCAT parameter on a DFDSS CONVERTV.

STGADMIN.ADR.CONVERTV
> Protects the ability to use DFDSS CONVERTV.

Authority to Activate an SMS Configuration

The FACILITY resource class profile named STGADMIN.IGD.ACTIVATE.CONFIGURATION protects the ability to activate SMS configurations. If you issue the ACTIVATE command from the Control Data Set Application Selection panel without FACILITY resource class authority, you remain on the panel without the ability to provide additional input until action is taken on the operator console. If you have authority to activate an SMS configuration, activation proceeds without confirmation from the operator console and you continue with normal operations.

FIELD Resource Class

The RACF FIELD resource class controls the ability of users to specify or update the following RACF profile fields:

- Resource owner, RESOWNER, for data set profiles
- Default SMS DATACLAS, STORCLAS, and MGMTCLAS values for user/group profiles
- Application identifier, DATAAPPL for user/group profiles.

Protecting SMS Storage and Management Classes

The *authorization* of SMS classes differs from the *protection* of SMS classes. For example, if you withhold authorization for a particular storage class, end users cannot specify that storage class in their JCL stream. End users specifying the storage class receive a JCL error.

Protection involves more than using an SMS class, it involves changing one. All end users might have READ access to a particular storage class, but only a select few will have UPDATE authority. Also, only people with UPDATE authority to SMS classes can update the SCDS.

Authorizing the Storage and Management Classes

You can authorize storage classes and management classes. Data classes do not require authorization.

RACF contains two new resource classes: STORCLAS and MGMTCLAS. You authorize storage classes and management classes by defining them as RACF profiles to the general resource classes. Note that the two new resource classes, storage class and management class, are distinct from SMS storage class and management class.

Protecting the Control Data Sets

You need to protect three types of SMS data sets:

1. Control data sets

   - SCDS
   - ACDS
   - COMMDS

2. ACS Routine Source Data Sets

   - Partitioned data sets whose members are source ACS routines
   - Sequential data sets that contain one source ACS routine

3. ACS Routine test library

   - Partitioned data sets whose members are ACS test cases.

You can issue SECURITY from the command line of the ISMF Data Set Application to authorize these data sets. Give read access to ACDS, COMMDS, and SCDS. Note that anyone who is authorized to update an SCDS or an ACS routine source data set is also authorized to change any items within these data sets.

Authorizing for TSO and ISPF

To define an SMS data set under TSO, you do not need proper Authorized Program Facility (APF) authorization, but it is recommended because APF authorization reduces path length. For the DEFINE and IMPORT commands, you can establish this authorization by adding the the DEFINE and IMPORT commands processor names to installation access list APFCTABL in CSECT IKJEFTE2. For a call of access method services, you can establish this authorization by adding IDCAMS to the installation access list APFPTABL in CSECT IKJEFTE8. To define or import an SMS data set while under option 6 of ISPF/PDF, you must include the DEFINE and IMPORT processor names in the command table in ISPTCM.

Storage Administrator Authorization

You should receive the following authority from the security administrator:

- CLAUTH to STORCLAS and MGMTCLAS resource classes

- Authority to activate an SMS configuration

- Authority to perform exceptional catalog and DFDSS operations on SMS-managed data sets

- Authority to all ISMF applications, dialogs, line operators, and commands

- Update authority to the RESOWNER field in the DATASET profiles

- Update authority to the DATAAPPL, MGMTCLAS, STORCLAS, and DATACLAS, fields in the USER/GROUP profiles.

It might be preferable to have a group profile with authorities mentioned above for all storage administrators.

Chapter 13. ACS Language Reference

ACS routines determine SMS classes and storage groups for all new data set allocations and for data set allocations that occur as a result of converting, copying, recalling, restoring, or moving data sets. You write ACS routines in the ACS language, which is a high-level programming language. You can write four ACS routines, one for each type of SMS class and storage group.

After writing the routines, you must translate them into an object form that SMS understands. A successful translation places the ACS object in a specified SCDS. After you activate the configuration contained in that SCDS, you can use the ACS routines to govern storage management.

When you enter the Automatic Class Selection Application of ISMF and select the EDIT option, you are linked to PDF Edit, where you can create or modify ACS routines. Leaving EDIT returns you to the Automatic Class Selection Application of ISMF, where you can translate, validate, or test any ACS routine.

This chapter contains four main sections. The first section describes the ACS language constants. The second section describes *read-write variables*. The third section describes *read-only variables*, which the ACS routines use for comparison operations. The fourth section describes the ACS language statements and illustrates the use of the statements in a storage group selection routine.

Constants

You can use four types of constants in ACS routines:

Numeric A numeric is a string consisting of up to ten characters, 0 through 9. You can use numerics in comparison operations involving the &NQUAL, &NVOL, and &RETPD read-only variables, which are discussed in "Read-Only Variables" on page 120.

KB, MB KB and MB are suffixes for numeric constants, such as 200KB and 10MB. You can use KB and MB in comparison operations involving the &SIZE and &MAXSIZE read-only variables, which are discussed in "Read-Only Variables" on page 120. The maximum prefix value for KB is 2,147,483,647. The maximum prefix value for MB is 2,097,151.

Literal A literal is a character string, such as 'SYS1.PARMLIB', that is enclosed in single quotation marks. The maximum length of a literal is 255 characters. If you want a literal to contain a single quotation mark, such as PAYROLL'SDATA, then you must specify two single quotation marks: 'PAYROLL''SDATA'.

Mask A mask is a character string, such as SYS1.*LIB, that is **not** enclosed in single quotation marks. You can use a mask to represent job names, volume serial numbers, or other system values that have a common string of characters, such as all volume serial numbers that begin with IMS. You can also use a mask to represent data set names that have a common string of characters.

A mask must begin with an alphanumeric character, numeric character (0-9), national character ($, @, #), asterisk (*), or percent sign

(%). The three characters "*", "%", and "." have special significance in a mask. The following sections describe the rules for using both the simpler masks and the slightly more involved data set masks. See "FILTLIST" on page 128 for an explanation of the use of the mask characters.

Simple Mask Rules

The following rules apply to the special characters in a simple mask consisting of a single level name:

- A "*" means that zero or more characters can be present.

- Two or more adjacent "*" are not allowed within a simple mask.

- A "%" is a single character place holder. "%%%" represents three character positions.

Simple Mask Examples

TSO*	All names of any length beginning "TSO"
XYZ	All names of any length having three adjacent characters "XYZ"
IMS%%%	All six-character names beginning "IMS"
***%WK%%**	All names where the second and third characters of the last five (or only five) are "WK"

Data Set Mask Rules

The following rules apply to the special characters in a data set mask:

- You can separate data set qualifiers with periods, ".".

- Each qualifier has a maximum length of eight characters. The maximum length for the entire data set mask is 44 characters.

- A "%" is a single character place holder. "%%%" represents three character positions.

- A "*" within a qualifier means that zero or more characters can be present.

- A single "*" can also represent a single qualifier place holder.

- A "**" means that zero or more qualifiers can be present.

- A "**" cannot appear with any other characters within a qualifier.

- Three or more adjacent "*" are not allowed within a qualifier.

Data Set Mask Examples

SYS1.**

All names where the first (or only) qualifier is "SYS1"

**.OUTLIST

All names where the last (or only) qualifier is "OUTLIST"

..PAYROLL.*.SALARY.*

All names with six qualifiers where the third qualifier is "PAYROLL" and the fifth qualifier is "SALARY"

.%%TEST..DATA

> All names with four qualifiers where the second qualifier has six characters ending in "TEST" and the fourth qualifier is "DATA"

.*ABC*.

> All names where some (or only) qualifier contains the characters "ABC" (or only "ABC")

LABMGR.**.DATA

> All names where the first qualifier is "LABMGR" and the last qualifier is "DATA"

Read-Write Variables

You write ACS routines to assign values to read-write variables. With the exception of STORGRP, you can also use read-write variables as values in comparison operations. The ACS language has three read-write variables:

&DATACLAS
&STORCLAS
&MGMTCLAS

If the read-write variables are explicitly specified by the user, they will have an initial value that may be overridden by the ACS routine. If the value is not overridden and the initial name is not defined in a currently active configuration, the allocation will fail and the value will be set to "?".

You must specify a read-write variable on the PROC statement of the corresponding ACS routine. See "PROC" on page 128 for details.

Each ACS routine can only set its corresponding read-write variable:

- The data class routine can only set &DATACLAS.
- The storage class routine can only set &STORCLAS.
- The management class routine can only set &MGMTCLAS.

The ACS routines assign values to read-write variables using the SET command, which is explained in "SET" on page 130.

In an ACS routine, you can assign an alphanumeric name enclosed in single quotation marks to the read-write variables. Also, you can assign a list of up to fifteen alphanumeric storage group names, each enclosed in single quotation marks, to the &STORGRP read-write variable.

Figure 72 indicates which read-write variables you can set, which ones you can use for comparisons, and which ones are invalid, for each of the ACS routines:

Figure 72. Using Read-Write Variables in ACS Routines

| ACS Routine | Read-Write Variable | | | |
	&STORGRP	&MGMTCLAS	&STORCLAS	&DATACLAS
Data class	Invalid	Compare	Compare	Set/Compare
Storage class	Invalid	Compare	Set/Compare	Compare
Management class	Invalid	Set/Compare	Compare	Compare
storage group	Set	Compare	Compare	Compare

Read-Only Variables

Most ACS variables are read-only. Read-only variables contain data set and system information, and they reflect what is known at the time of the allocation request. You can use read-only variables in comparison operations, but you cannot change their values.

All of the read-only variables appear below in Figure 73. The following pages explain the uses of the read-only variables.

Figure 73. Read-Only Variables

&ACCT_JOB	&DD	&DSOWNER	&LLQ	&RECORG
&ACCT_STEP	&DEF_DATACLAS	&DSTYPE	&MAXSIZE	&RETPD
&ACSENVIR	&DEF_MGMTCLAS	&EXPDT	&MSVGP	&SIZE
&ALLVOL	&DEF_STORCLAS	&GROUP	&NQUAL	&UNIT
&ANYVOL	&DSN	&HLQ	&NVOL	&USER
&APPLIC	&DSORG	&JOB	&PGM	&XMODE

Name	Description
&ACCT_JOB	The accounting information from the JOB card. (For a description of the indexing function for accounting information, see "Special Functions" on page 126.)
	Type: Literal
	Max value: 142 characters
&ACCT_STEP	The accounting information from the EXEC card. This information will be refreshed for each step in the job. (For a description of the indexing function for accounting information, see "Special Functions" on page 126.)
	Type: Literal
	Max value: 142 characters
&ACSENVIR	The environment in which the ACS routine was invoked, one of:

'RECALL'   for data set recall operations
'RECOVER'   for data set recover operations
'CONVERT'   for data set convert in place operations

'ALLOC'	for new data set allocations (this is the default)
'other'	installation exit can set its own value before re-invoking ACS

Type: Literal
Max value: 8 characters

&ALLVOL   The volume serial numbers specified for data set allocation when &ACSENVIR is not recall or recover. When the environment is recall or recover, &ALLVOL is either the volume serial number on which the data set resided at the time it was migrated or backed up, or the volume serial number specified as the target volume of the recall or recover. See 126 for usage information.

Type: Literal
Max value: 6 characters

&ANYVOL   The volume serial numbers explicitly specified for the volume(s) if &ACSENVIR is not recall or recover. If the environment is recall or recover, &ANYVOL is either the serial number on which the data set resided at the time it was migrated or backed up, or the volume serial number specified as the target volume of the recall or recover. See 126 for usage information.

Type: Literal
Max value: 6 characters

&APPLIC   The name of the application associated with the resource owner of the data set (set only if RACF installed and ACSDEFAULTS is YES in IGDSMSmm).

Type: Literal
Max value: 8 characters

&DD   DD (data definition) statement describing the data set.

Type: Literal
Max value: 8 characters

&DEF_DATACLAS   The data class name associated with the resource owner of the data set (set only if RACF is installed and ACSDEFAULTS is YES in IGDSMSmm).

Type: Literal
Max value: 8 characters

&DEF_MGMTCLAS   The management class name associated with the resource owner of the data set (set only if RACF is installed and ACSDEFAULTS is YES in IGDSMSmm).

Type: Literal
Max value: 8 characters

&DEF_STORCLAS   The storage class name associated with the resource owner of the data set (set only if RACF is installed and ACSDEFAULTS is YES in IGDSMSmm).

Type: Literal
Max value: 8 characters

&DSN

The name of the data set for which ACS processing is taking place. For VSAM data sets, only the cluster name is passed to the ACS routine; the component name(s) are not.

Type: Literal
Max value: 44 characters

&DSORG

The data set organization, one of:

'PS'	Physical sequential
'PO'	Partitioned
'VS'	VSAM organization
'DA'	BDAM organization
'null'	Indeterminate (for example, it can be specified in a program DCB)

Type: Literal
Max value: 2 characters

&DSOWNER

The name of the user or group that owns the data set (set only if RACF is installed).

Type: Literal
Max value: 8 characters

&DSTYPE

The data set type, one of:

GDS	One generation data set of a generation data group
PERM	Standard permanent data sets
TEMP	Temporary data sets
null	Indeterminate

Type: Literal
Max value: 8 characters

&EXPDT

The expiration date (in the form, 'YYYYDDD').

Type: Literal
Max value: 7 characters

&GROUP

The name specified on the GROUP= JCL keyword. If the environment is recall or recover, &GROUP is set only if the requester of the recall or recover is not a DFHSM authorized user. When DFHSM invokes the ACS routines, &GROUP is the group associated with &USER.

Type: Literal
Max value: 8 characters

&HLQ

The high-level (first) qualifier of the data set name.

Type: Literal
Max value: 8 characters

&JOB

The job name, the started task name, or the TSO/E userid from the JOB card, depending on your execution mode (&XMODE). When DFSM invokes the ACS routines, &GROUP is the group associated with &USER.

Type: Literal
Max value: 8 characters

&LLQ	The low-level (last) qualifier of the data set name.
	Type: Literal
	Max value: 8 characters
&MAXSIZE	The maximum size (in KB or MB) of a new data set. For non-VSAM data sets, the value is &SIZE plus 15 extents. For VSAM data sets, the value is &SIZE plus 123 extents.
	Type: Numeric
	Max value: 2,147,483,647 for KB Max value: 2,097,151 for MB
&MSVGP	The group name for mass storage (MSS) volumes.
	Type: Literal
	Max value: 8 characters
&NQUAL	The number of qualifiers in the data set name.
	Type: Numeric
	Max value: 22
&NVOL	The maximum of the volume count, UNIT count, and sum of explicit VOL=SER specifications.
	Type: Numeric
	Max value: 2,147,483,647
&PGM	The name of the program the system is running.
	Type: Literal
	Max value: 8 characters
&RECORG	The data set record organization, one of:

&LLQ — The low-level (last) qualifier of the data set name.

Type: Literal
Max value: 8 characters

&MAXSIZE — The maximum size (in KB or MB) of a new data set. For non-VSAM data sets, the value is &SIZE plus 15 extents. For VSAM data sets, the value is &SIZE plus 123 extents.

Type: Numeric
Max value: 2,147,483,647 for KB Max value: 2,097,151 for MB

&MSVGP — The group name for mass storage (MSS) volumes.

Type: Literal
Max value: 8 characters

&NQUAL — The number of qualifiers in the data set name.

Type: Numeric
Max value: 22

&NVOL — The maximum of the volume count, UNIT count, and sum of explicit VOL=SER specifications.

Type: Numeric
Max value: 2,147,483,647

&PGM — The name of the program the system is running.

Type: Literal
Max value: 8 characters

&RECORG — The data set record organization, one of:

'KS'	VSAM key sequenced (KSDS)
'ES'	VSAM entry sequenced (ESDS)
'RR'	VSAM relative record (RRDS)
'LS'	VSAM linear
'null'	SAM PS or PO data set

Type: Literal
Max value: 2 characters

&RETPD — The retention period (nnnn days).

Type: Numeric
Max value: 2,147,483,647

&SIZE — The primary amount of space (in KB or MB) requested for a new data set on a DASD volume.

Type: Suffixed numeric
Max value: 2,147,483,647 for KB
Max value: 2,097,151 for MB

&UNIT — IBM-supplied or installation defined generic name for a device type (for example 3380, SYSDA).

Type: Literal
Max value: 6 characters

&USER	The userid of the person creating the data set. When DFHSM invokes the ACS routines, &USER is either the requestor of the recall or recover, or the userid of the DFHSM task. If the environment is recall or recover, the variable is set only if the requestor of the recall or recover is not a DFHSM authorized user.

Type: Literal
Max value: 8 characters

&XMODE	The execution mode in which the data set is being created, one of:

'BATCH' Batch execution mode
'TSO' TSO execution mode
'TASK' A started task

Type: Literal
Max value: 8 characters

Usage Notes for Read-Only Variables

SMS derives the values of read-only variables before it invokes the ACS routines. The values are based on what is known at the time of the allocation request (for example, if a unit name has been specified on the allocation request, then the &UNIT variable contains the specification).

R/O variables are defaulted from DATA CLASS (if assigned) before SC, MC, and SG ACS routines are invoked.

For a partitioned data set, &DSN will only consist of the library name. A member name, if specified, will not be part of the value of the &DSN variable. There is no class selection capability based on member name.

Not all read-only variable values are meaningful during the actual operation of any ACS routine. For example, the value of &USER could be null in the storage group selection routine if no userid had been specified on the JOB card or determined from the environment.

Constraints on Read-Only Variable Usage

The following read-only variables are **not** allowed in the storage group selection routine:

&ACCT_JOB
&ACCT_STEP
&DD
&JOB
&MSVGP
&PGM
&XMODE
&USER

The following read-only variables are **not** passed to any ACS routine for other than initial allocation (&ACSENVIR = ALLOC):

&ACCT_JOB
&ACCT_STEP
&DD

&JOB
&MSGVP
&XMODE

Comparison Operators

Comparison operators allow you to determine the difference between two values. The following comparison operators are allowed:

This	Means this
GT or >	Greater than
LT or <	Less than
NG or ¬ >	Not greater than
NL or ¬ <	Not less than
EQ or =	Equal
NE or ¬ =	Not equal
GE or > =	Greater than or equal
LE or < =	Less than or equal

Alphabetic characters are sorted before digits (A-Z come before 0-9). The following comparison is true for all high-level qualifiers alphanumerically greater than "M":

```
IF &HLQ > 'M' THEN . . .
```

For FILTLIST or mask comparisons, only EQ and NE are valid. See "FILTLIST" on page 128 for details.

Comparison Rules

The following rules apply to comparisons:

1. For a comparison to be valid, one operand must either be a read-only variable or a read-write variable and the other operand must be a literal (any of the four types), a read-only variable, or a FILTLIST name.

2. Numerics are right justified.

3. Literals are left justified and padded with blanks.

4. Type checking is done to ensure that numeric read-only variables are not being compared to characters (literals) and that character (literal) read-only variables are not being compared to numbers. &NQUAL, &NVOL, &SIZE, &MAXSIZE and &RETPD are the only numeric read-only variables.

5. Limited length checking of read-only variables with their maximum length values is performed to ensure that the maximum lengths are not exceeded. For example, the literal to which &DSN is being compared must be no longer than 44 characters.

   See "Read-Only Variables" on page 120 for maximum lengths.

Boolean Expressions

You can use the following Boolean operators in any ACS routine:

This	Means this
AND or &&	And
OR or |	Or

AND takes precedence over OR. In the following example, condition 1 and condition 2, together with the AND operator, are evaluated first; their result is then evaluated with condition 3 and the OR operator.

```
IF  (condition 1) AND
    (condition 2) OR
    (condition 3) THEN
```

You can group relational expressions using sets of parentheses. You can combine these groupings with the Boolean operators to form more complicated relational expressions. In the following example, each set of expressions within parentheses is evaluated first; the two results are then evaluated.

```
IF (condition 1 AND condition 2) AND
   (condition 3 OR condition 4) THEN . . .
```

Special Functions

The ACS language provides the following three indexing functions, which allow you to make class selections based on specific details about each data set.

Selection Routine	Data Set Qualifier	&ALLVOL &ANYVOL	Accounting Information
Data class	Yes	Yes	Yes
Storage class	Yes	Yes	Yes
Management class	Yes	Yes	Yes
Storage group	Yes	Yes	No

Data Set Qualifier Indexing Function: Allows you to index the &DSN variable (for accessing particular data set qualifiers):

&DSN(1)
> first qualifier of the data set name

&DSN(3)
> third qualifier of the data set name

&DSN(&NQUAL)
> last qualifier of the data set name

The only accepted values for indexes are numbers (1 through 22) and the read-only variable, &NQUAL.

&ALLVOL and &ANYVOL Function: Allow you to compare the volume serial number(s) explicitly specified on input with a comparison variable (for example, a FILTLIST variable). The use of &ALLVOL in a comparison expression returns a true value if ALL of the input volsers satisfy the desired condition. The use of

&ANYVOL will return a true value if ANY of the input volsers satisfies the desired condition. For example, let IMS101, IMS102, and TSO191 be the input volsers to the following routine:

```
PROC STORCLAS

    IF &ALLVOL = IMS* THEN
        (stuff left out)
    SELECT
      WHEN (&ANYVOL = TSO*)
        (stuff left out)
    END
END
```

The IF statement is false, because *all* volumes do not match. While IMS101 and IMS102 satisfy the TSO* mask, TSO191 does not.

The WHEN statement is true, because *any* (at least one) volume does match. While IMS101 and IMS102 fail to satisfy the IMS* mask, TSO191 does satisfy the mask.

Note: For a detailed explanation of the IF statement, see "IF" on page 131.

Accounting Information Indexing Function: Allows you to reference specific fields in the JOB or STEP account information. A field is defined as a unit of data which is separated by comma(s) in the account information. The use of &ACCT_JOB(n) and &ACCT_STEP(n), where n is the field number, indicates that indexing is requested. If either &ACCT_JOB or &ACCT_STEP is specified without an index, then the default action is to access the first field of the accounting information (for example, &ACCT_JOB(1) or &ACCT_STEP(1) will result by default).

```
&ACCT_JOB(1)     first field of the JOB accounting information
&ACCT_JOB        first field of the JOB accounting information (default)
&ACCT_JOB(3)     third field of the JOB accounting information
&ACCT_STEP       first field of the STEP accounting information (default)
&ACCT_STEP(5)    fifth field of the STEP accounting information
```

The only accepted values for indexes are numbers from 1 to 71.

Statements

This section describes the function and syntax of the ACS language statements that you can use when writing ACS routines.

The continuation characters "+" and "-" allow you to extend literal constants to the next line. To ignore the leading blanks on the following line, use "+". If you want to include the leading blanks on the next line as part of a literal, then use a "-". You cannot continue masks, numbers, KB or MB numerics, or keywords.

The maximum number of nesting levels for any combination of ACS statement types is thirty-two. (For example, a nested IF statement is one that appears within an IF statement.)

Comments begin with a slash-asterisk pair, "/*", and the end with an asterisk-slash pair, "*/". You can place comments anywhere within an ACS routine.

The statement types are defined as follows:

Statement	Description
PROC	Start of an ACS routine.
FILTLIST	Definition of filter criteria.
SET	Assigns a value to a Read/Write variable.
DO	Start of statement group.
IF	Provides conditional statement execution.
SELECT	Defines a set of conditional execution statements.
EXIT	Causes immediate termination of the ACS routine and can be used to force allocation failures.
WRITE	Sends messages to the end user.
END	End of statement group (for example, DO, SELECT) or ACS routine (PROC).

PROC

PROC is the first statement of each ACS routine. It identifies the ACS routine and which read-write variable the routine sets. You can precede the PROC statement with blank lines or comments, but not with other statements. For TSO CLIST compatibility, you can place a blank and then a number, such as 0 or 1, after the PROC statement. The number does not affect ACS language processing. To identify an ACS routine and the value it is to determine, you must specify a read-write variable at the end of the PROC statement. You must also place an END statement at the end of each ACS routine.

Syntax

PROC $<n>$ *read-write variable*

END

- *n* is optional and can contain any numeric value.

- *read-write variable* is a mandatory value that can be DATACLAS, STORCLAS, MGMTCLAS, or STORGRP. You can optionally precede the variable with an ampersand, &.

Examples

```
PROC 1 DATACLAS
PROC 0 &STORCLAS
PROC &MGMTCLAS
PROC STORGRP
```

FILTLIST

The FILTLIST statement is a definition list that you can use when testing variables in an ACS routine. You define the information that you want to include and exclude in the list using the INCLUDE and EXCLUDE keywords. Then you can compare read-only variables to items in the list using IF-THEN and SELECT-WHEN statements, without having to write elaborate AND and OR combinations.

FILTLIST is a definition statement that simplifies comparison operations. It is not an execution statement, and it does not change the value of any variables.

Because a FILTLIST can contain only literal values, you can only compare it to literal read-only variables. This excludes the numerically valued &NQUAL, &NVOL, &SIZE, &MAXSIZE, and &RETPD read-only variables from FILTLIST comparisons.

You must define a FILTLIST before you reference it in the body of an ACS routine.

Syntax

FILTLIST *name* [< INCLUDE(*list*) >] [< EXCLUDE(*list*) >]

- *name* is mandatory and can be up to 31 alphanumeric characters in length. You can also use an underscore, _, but it cannot be the first character. In the FILTLIST, you can optionally precede the name with an ampersand, &. When referring to the FILTLIST in the body of the routine, you must always precede the FILTLIST name with an ampersand.

- You must specify INCLUDE, EXCLUDE, or both in the FILTLIST statement. If a list item satisfies both the INCLUDE and EXCLUDE criteria, EXCLUDE takes precedence and prevents the item from being included in the list.

- *list* can contain literals, simple masks, and data set masks.

Example

```
PROC STORCLAS

    FILTLIST VLIST2 INCLUDE(DBX*, TSO*) EXCLUDE('DBX191', 'TSO256')

    IF &ALLVOL = &VLIST2 THEN
        (some action)

END
```
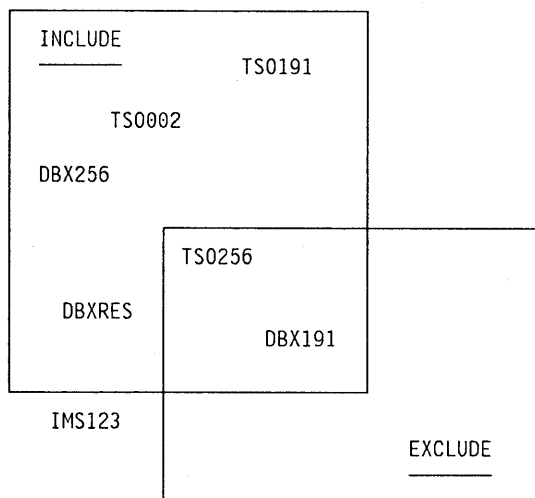

Figure 74. Using the INCLUDE and EXCLUDE

In the environment shown in Figure 74, the value of the IF statement is *true* for any of the following volume serials:

TSO191
TSO002
DBX256
DBXRES

The value of the IF statement is *false* for the following volume serials, because they match the EXCLUDE filter criteria:

TSO256
DBX191

SET

The SET statement assigns values to the read-write variables. The values can be the names of constructs or respective RACF derived defaults (&DEF_DATACLAS, &DEF_STORGRP, &DEF_STORCLAS, or &DEF_MGMTCLAS). You can assign one name to &DATACLAS, one name to &STORCLAS, and one name to &MGMTCLAS, but you can assign a list of up to 15 names to &STORGRP.

The names can be from one to eight characters long, and they must be enclosed in single quotation marks. The individual names belonging to a &STORGRP must be enclosed in single quotation marks and separated by commas. A name must begin with either an alphabetic or national ($, @, #) character, and the remaining characters can be alphabetic, numeric, or national.

In the body of an ACS routine, you can only set the value of the read-write variable identified in the PROC statement. You cannot set a read-write variable equal to another read-write variable, a FILTLIST variable, or a read-only variable (except for the RACF-derived values). For example,

```
SET &STORCLAS = &PGM
```

is invalid.

You can assign a null value to any of the read-write variables except for &STORGRP. You can assign a null value by specifying two single quotation marks with nothing between them (' ').

Syntax

SET *read-write variable* = *value*

- *read-write variable* is a mandatory value that can be &DATACLAS, &STORCLAS, &MGMTCLAS, or &STORGRP.

- You can specify EQ in place of the equals sign, =.

- *value* can be one name, a null value, a RACF read-only variable name, or a list of names in the case of the storage group ACS routine. Figure 75 summarizes the possible assignments of *value*.

Figure 75. Read-Write Variable Assignments				
	One Name in Single Quotation Marks	Null Value	RACF Read-Only Variable Name	List of Quoted Names Separated by Commas
Storage group	x			x
Management class	x	x	x	
Storage class	x	x	x	
Data class	x	x	x	

Examples

```
SET &STORCLAS EQ 'SCNORM'
SET &STORCLAS = 'SCNORM'
SET &DATACLAS = &DEF_DATACLAS          /* RACF read-only variable */
SET &MGMTCLAS = 'SCRATCH5'
SET &STORGRP = 'SG1','SG2','SG3'       /* list of values          */
SET &DATACLAS = ''                     /* null value assignment   */
```

Restrictions

You must have RACF installed and ACSDEFAULT(YES) specified in the IGDSMSnn to assign the &DEF_DATACLAS, &DEF_STORCLAS, or &DEF_MGMTCLAS otherwise a null value will be assigned to a read-write variable.

DO

You can group a collection of ACS language statements using a DO statement paired with an END statement. The DO statement can follow an IF-THEN clause, an ELSE clause, or a SELECT-WHEN group.

Syntax

DO
 <group of statements>
END

- *group of statements* can consist of zero or more ACS language statements.

Example

```
IF &HLQ='PAYROLL' THEN
DO
  WRITE 'No Payroll allowed'
  EXIT CODE(1)
END
```

IF

Use the IF statement for conditional statement execution. You must always follow an IF statement with a THEN clause. The THEN clause can be a single statement or a DO-END group of statements. You can optionally follow a THEN clause with an ELSE clause. If you specify the ELSE clause, you must follow it with either a single statement or a DO-END group of statements. If you want to specify the ELSE clause but do not want to follow it with an executable statement, follow it with an empty DO-END pair.

If the result of the IF statement comparison is *true*, the THEN clause is executed. If the result is false, the ELSE clause is executed. If you omit the ELSE, processing continues with the next sequential statement after the THEN clause.

Syntax

IF *relational expression* THEN *clause* <ELSE *clause*>

- *relational expression* can be a single comparison or it can be multiple comparisons joined by Boolean operators.

- THEN is a mandatory keyword.

- *clause* can contain a single statement, a DO-END group of statements, or a SELECT statement.

- ELSE is an optional keyword.

- An ELSE *clause* is mandatory if you specify the ELSE keyword.

Note: For information on what constitutes a valid relational expression, see "Comparison Rules" on page 125.

Example

Example of a specified ELSE:

```
IF &DSOWNER = 'BILL' THEN
  SET &STORCLAS = 'ONE'
ELSE
  SET &STORCLAS = 'TWO'
```

Example of a null ELSE:

```
IF &DSOWNER = 'BILL' THEN
    SET &STORCLAS = 'ONE'
ELSE
    DO END
SET &STORCLAS = 'TWO'
```

The statement on the left sets &STORCLAS equal to ONE or TWO, depending on the value of &DSOWNER.

The statement on the right set &STORCLAS equal to TWO, regardless of the value of &DSOWNER. If &DSOWNER equals 'BILL', then &STORCLAS is set to 'ONE'. The ELSE clause is skipped, and execution falls to the next statement, which changes &STORCLAS to 'TWO'. If &DSOWNER does not equal 'BILL', execution falls to the ELSE, which results in no assignment to &STORCLAS. Then execution proceeds to the next sequential statement, which sets &STORCLAS to 'TWO'.

SELECT

Use the SELECT statement to write conditional statements in sequential form rather than IF-THEN-ELSE form. A SELECT statement consists of a SELECT keyword, one or more WHEN clauses, an optional OTHERWISE clause, and an END statement. You can specify the SELECT statement in one of two forms. In the first form shown below, you include the variable being tested after the SELECT statement. In the second form shown below, you include the variable being tested after the WHEN keyword.

The first true WHEN condition is executed, and the remaining WHEN conditions are ignored.

Syntax (first form)

SELECT *(variable)*
WHEN *(value)* *<action>*

```
        .
        .
        .
```

*<*WHEN *(value)* *<action>* *>*
*<*WHEN *(value)* *<action>* *>*
*<*OTHERWISE *<action>* *>*
END

- *variable* can be any read-only or read-write variable.

- At least one WHEN keyword is mandatory.

- *value* can be a constant or a FILTLIST name.

- *action* can be an ACS statement, SELECT group, or a DO END group.

- OTHERWISE is an optional keyword.

Example (first form)

```
SELECT (&DSOWNER)
  WHEN ('IBMUSER1') SET &STORCLAS = 'PAYROLL'
  WHEN ('IBMUSER2') SET &STORCLAS = 'TEST'
  WHEN ('IBMUSER3') SET &STORCLAS = 'DEVELOP'
  OTHERWISE         SET &STORCLAS = 'NORMAL'
END
```

Syntax (second form)

SELECT
WHEN *(relational expression)* *<action>*

```
        .
        .
        .
```

*<*WHEN *(relational expression)* *<action>* *>*
*<*WHEN *(relational expression)* *<action>* *>*

*<*OTHERWISE *<action>* *>*
END

- At least one WHEN keyword is mandatory.

- *relational expression* can be a single comparison or it can be multiple comparisons joined by Boolean operators.

- *action* can be an ACS statement, SELECT group, or a DO END group.

- OTHERWISE is an optional keyword.

Example (second form)

```
SELECT
  WHEN (&DSOWNER = 'IBMUSER') SET &STORCLAS = 'PAYROLL'
  WHEN (&DSOWNER = 'IBMUSER2')

    IF &ACCT_JOB = '1234' THEN
      SET &STORCLAS = 'TEST'
    ELSE
      SET &STORCLAS = 'EVERYONE'

  WHEN (&DSTYPE = 'UNMOVE')  SET &STORCLAS = ''
  WHEN (&HLQ    = 'CADAM')   SET &STORCLAS = ''
  OTHERWISE                  SET &STORCLAS = 'COMMON'
END
```

EXIT

The EXIT statement immediately terminates the operation of an ACS routine.

Syntax

EXIT <CODE(*n*)>

- CODE is an optional keyword.

- *n* is an exit code. A nonzero value for *n* causes the subsequent ACS routines to be skipped and the allocation to fail with no explicit value assigned to the read-write variable in the ACS routine. If you do not specify a value for *n*, it assumes the default value of zero.

Examples

```
PROC STORCLAS
   FILTLIST SECVOL INCLUDE(PAY*, REC*) EXCLUDE('PAYR20', 'REC195')
   FILTLIST VALID_UNITS  INCLUDE('3330','3330V','3340','3350','3375',
                                                '3380','SYSDA','')

   IF &UNIT ¬= &VALID_UNITS THEN

      DO
         SET &STORCLAS = ''
         EXIT
      END

   IF &ALLVOL ¬= &SECVOL THEN EXIT CODE(22)

END
```

If the first IF statement is *true* (&UNIT does not match any unit named in the VALID_UNITS filter criteria), then execution of this ACS routine terminates immediately. Allocation proceeds, because the exit code is zero, the default.

If the second IF statement is *true* (none of the input volumes match the SECVOL FILTLIST criteria), then execution of this ACS routine terminates immediately and the allocation fails. The value for CODE, 22, is set and displayed as part of the 'allocation failed' error message written to the end user.

WRITE

Use the WRITE statement to issue a message to an end user at execution/allocation time. With the WRITE statement, you can notify end users that your are removing a particular storage class, you can inform end users that they lack sufficient authority to use a particular management class, or you can tell an end user that you have moved a tape data set to DASD.

In a TSO/E ALLOCATION environment, the text of the WRITE message displays only if the allocation fails.

A WRITE message can contain up to 110 characters of text and variables. The message substitutes the value of a variable for the following variable names:

&DATACLAS

&STORCLAS

&MGMTCLAS

&DEF_DATACLAS

&DEF_STORCLAS

&DEF_MGMTCLAS

You must enclose the message in single quotation marks. If you want a single quotation mark to be part of the message, use two single quotation marks to represent it.

This:	Appears as this:
WRITE 'This line''s short.'	This line's short.

You can use continuation characters (+, -) to continue text onto a subsequent line. The closing single quotation mark signifies the end of text.

A nine-character system message id and a single blank character precede your message to the end user. At execution/allocation time, an end user can receive a maximum of five messages. If any more messages are generated, a sixth and final message indicates that additional messages have been generated, but the additional messages are not displayed.

Syntax

WRITE '*message*'

- *message* is passed back to the end user.

Example

Assuming the value of &STORCLAS is 'SC1:', the following WRITE message

```
WRITE 'WARNING - &STORCLAS SPECIFIED (' &STORCLAS ') IS NOT ALLOWED'
```

Displays as:

```
IGD01005I WARNING - &STORCLAS SPECIFIED (SC1) IS NOT ALLOWED
```

END

The END statement concludes an ACS routine, a DO group, and a SELECT statement.

Syntax

END

Examples

```
PROC STORCLAS
 (source code)
END
```

Sample ACS Routine

The following example illustrates some reasonable techniques for properly using the ACS routines.

```
PROC STORCLAS

/***********************************************************************/
/*                                                                     */
/* THIS IS THE PRODUCTION SELECTION SPECIFICATION FOR SETTING STORCLAS */
/*                                                                     */
/***********************************************************************/

FILTLIST DBVOLS  INCLUDE(IMS*,DB2*)                   /* ALL DATA BASE VOLUMES */
                 EXCLUDE('IMS053','DB2007')

FILTLIST DBJOBS  INCLUDE(IMS*,PROD*,ACCT*)            /* ALL DATA BASE JOBS */

FILTLIST VALID_UNITS INCLUDE('3330','3330V','3340','3350','3375',
                      '3380','SYSDA')                 /* VALID UNITS FOR SMS */

IF &UNIT ¬= &VALID_UNITS

    THEN DO
            SET &STORCLAS = ''
            WRITE 'INVALID UNIT TYPE FOR SMS ALLOCATION'
            EXIT
         END

SELECT

        WHEN &DSN = SYS1.**                              /* SYSTEM DATA */
            SET &STORCLAS = 'SYSTEM'

        WHEN (&ALLVOL = &DBVOLS) && (&JOB = &DBJOBS)  /* DATA BASE DATA */
            SET &STORCLAS = 'DBPOOL'

        WHEN (&DSN(3) = 'CLEAR') | (&ANYVOL ¬= TSO*)   /* NON-SMS DATA */
            SET &STORCLAS = ''

        WHEN &DEF_STORCLAS ¬= ''                        /* IF DEFAULTS EXIST */
            SET &STORCLAS = &DEF_STORCLAS

        OTHERWISE SET &STORCLAS = 'COMMON'              /* ALL OTHER DATA */

END                                                     /* END STORCLAS PROC */
```

Figure 76. Production ACS Routine for Storage Class
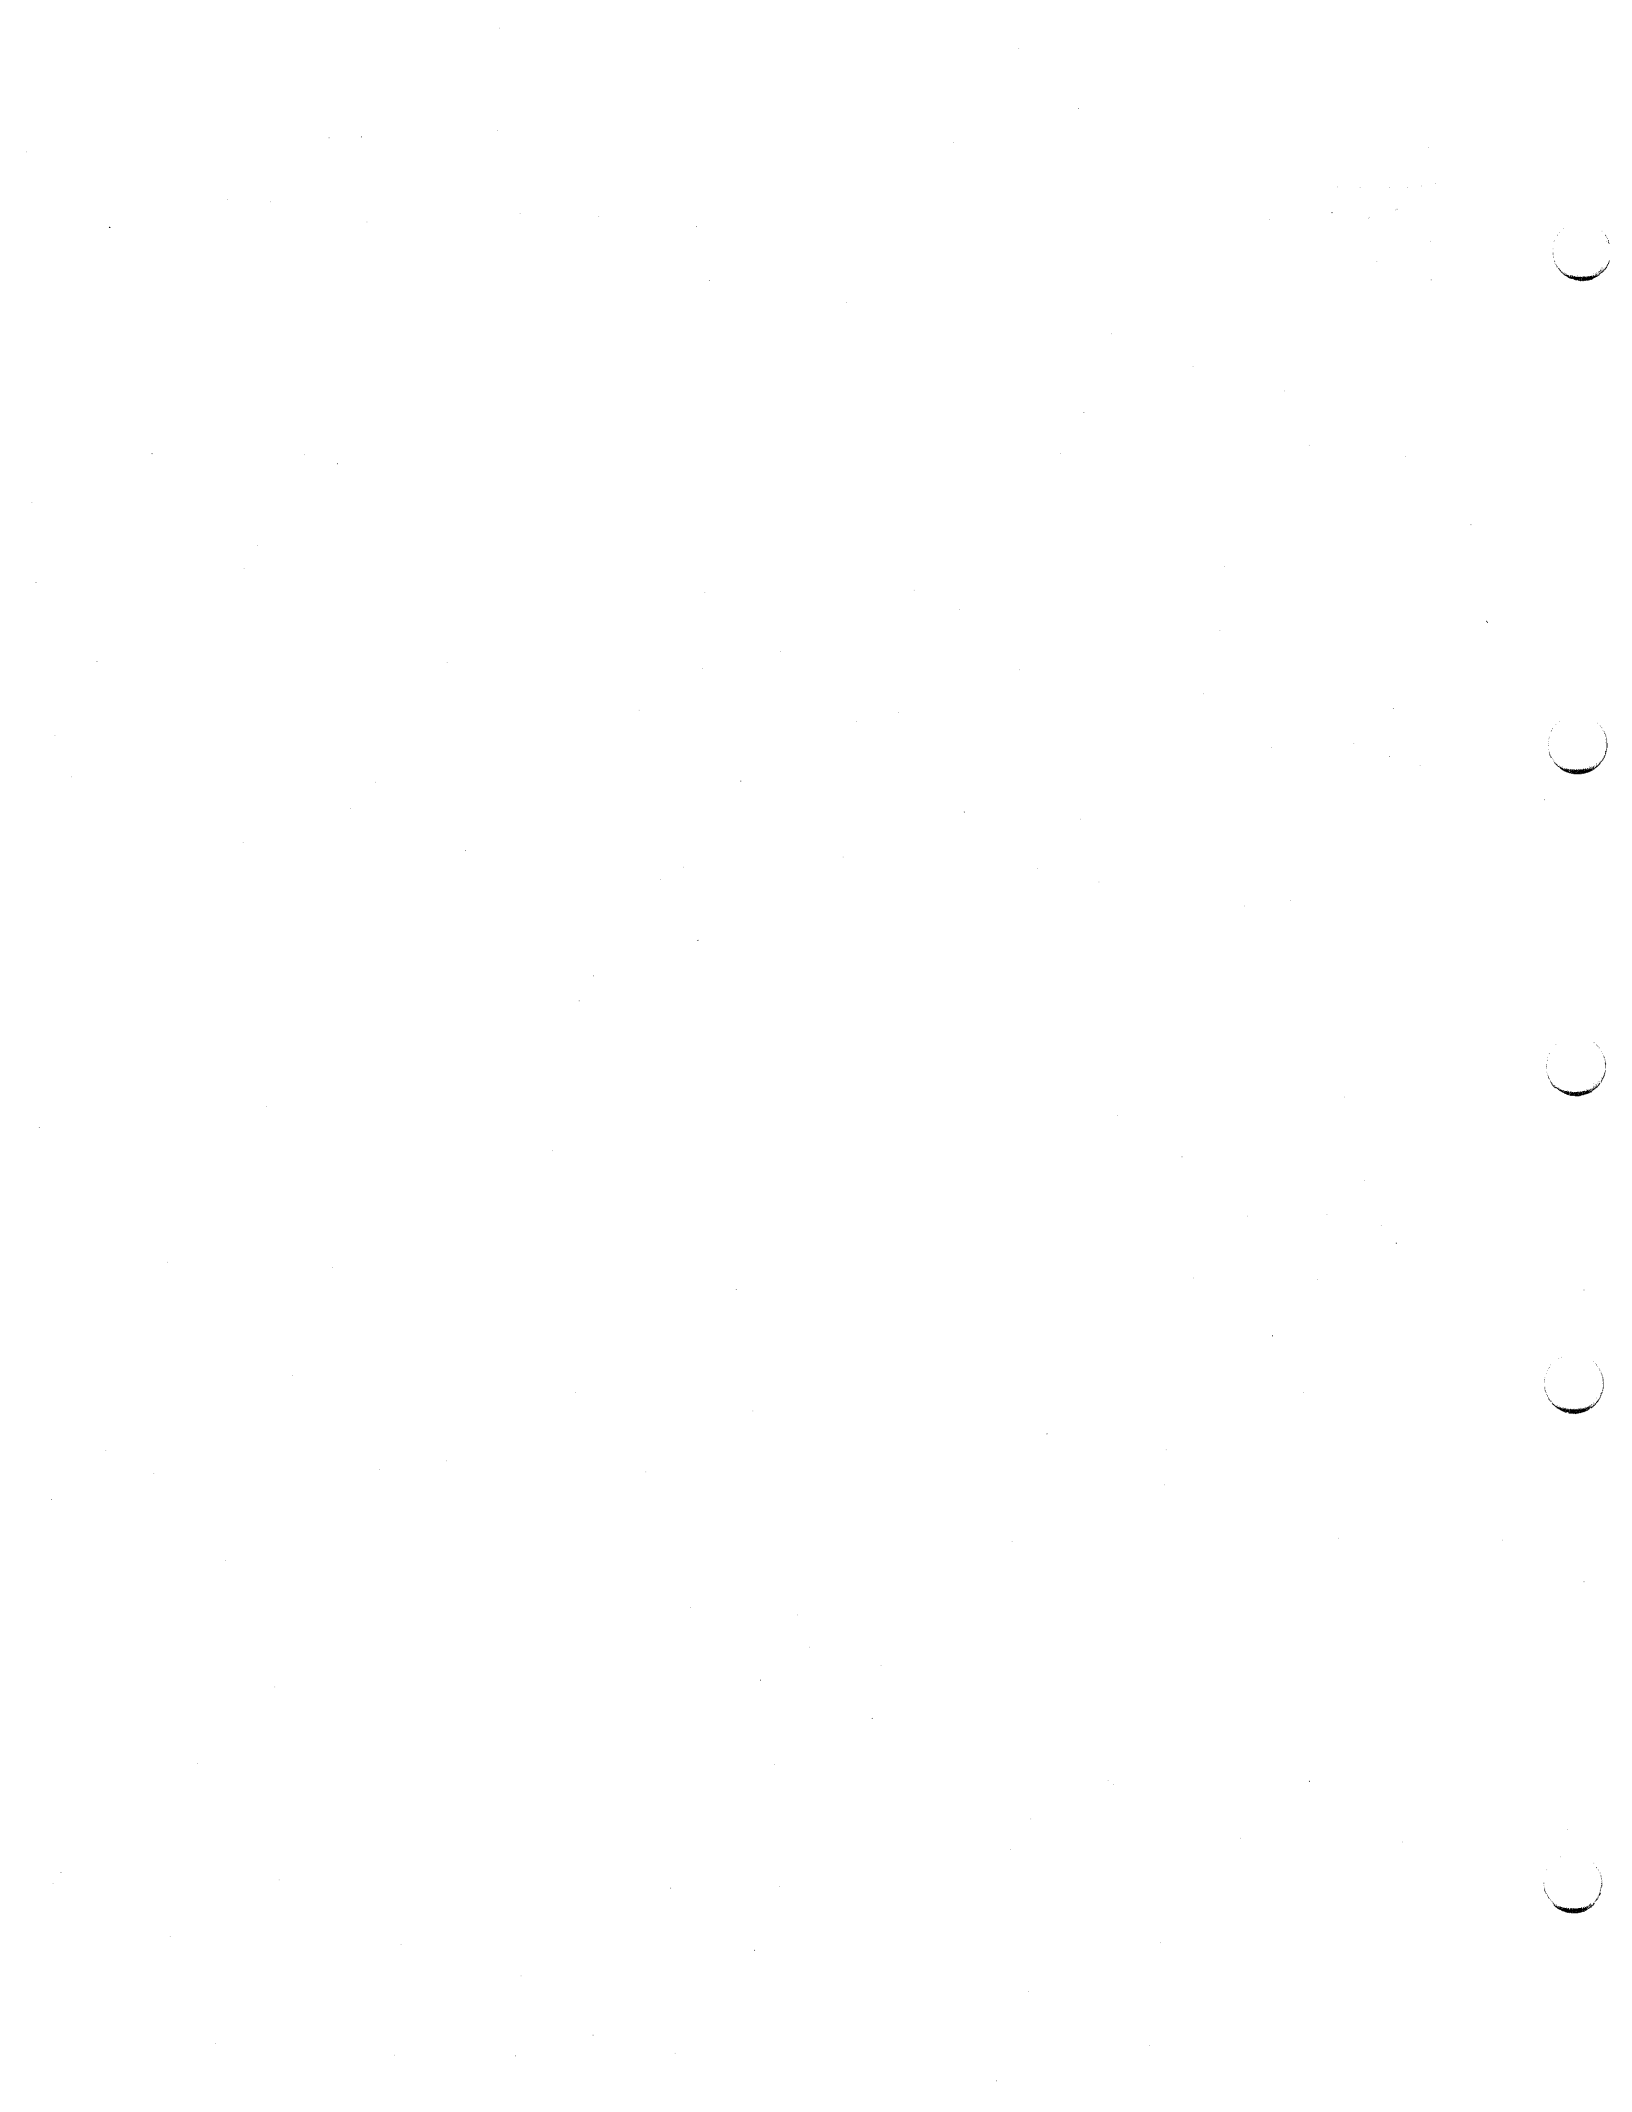
Appendix A. ISMF Command and Line Operator Reference Summary

The abbreviations in the Application column of these tables represent the following:

DS	Data set
VOL	Volume
DC	Data class
SC	Storage class
MC	Management class
SG	Storage group
CDS	Control data set
ACS	Automatic class selection

Note: If you specify an equal sign after any DFDSS or DFHSM line operator, processing occurs in last-use mode and MVS/DFP does not display an entry panel.

The following table lists the ISMF commands and line operators available to storage administrators. See the *MVS/ESA Interactive Storage Management Facility User's Guide* for the line operators and commands available to the end user.

Figure 77 (Page 1 of 2). ISMF Commands

Command	Minimum Abbreviation	Description	Application	Source
ACTIVATE	AC	Copy the contents of an SCDS into an ACDS and activate it, or activate an existing ACDS.	CDS	MVS/DFP
CANCEL	CA	Return to the previous dialog without performing any of the current dialog functions.	DC, SC, MC, SG, ACS, CDS	MVS/DFP
CLEAR	CLE	Reset line operator history.	DS, VOL, DC, SC, MC, SG	MVS/DFP
CLEAR PAGEx	CL PAGEx	Clear pages on selection entry panels and filter panels.	DS, VOL,	MVS/DFP
CLEAR PAGE	CL PA	Clear pages on selection entry panels and filter panels.	DS, VOL,	MVS/DFP
COMPRESS	COM	Reclaim embedded unused space from a list of PDSs.	DS	DFDSS
COPY	COP	Copy a list of data sets to a DASD volume of like or unlike device type.	DS	DFDSS
DOWN	DO	Scroll forward the specified number of list entries.	DS, VOL, DC, MC, SC, SG	MVS/DFP
DUMP	DU	Dump a list of data sets to tape, DASD, or mass store volumes.	DS	DFDSS
END	END	Exit the current ISMF function or panel and return to the previous panel.	All	MVS/DFP
ERTB	ER	Display the ISMF Error Table.	All	MVS/DFP

Figure 77 (Page 2 of 2). ISMF Commands

Command	Minimum Abbreviation	Description	Application	Source
FILTER	FIL	Tailor the list to include only specific entries.	DS, VOL, MC, DC, SC, SG	MVS/DFP
FILTER CLEAR	FIL C	Clear the filter entries but bypass the entry panel.	DS, VOL	MVS/DFP
FIND	FIN	Find a specific data column.	DS, VOL, MC, DC, SC, SG	MVS/DFP
FOLD	FO	Extend the data set name column.	DS	MVS/DFP
HELP	HELP	Request the help panels associated with the current panel.	All	MVS/DFP
LEFT	L	Scroll left the specified number of columns.	DS, VOL, DC, SC, MC, SG	MVS/DFP
PROFILE	P	Invoke the ISMF profile.	All	MVS/DFP
RELEASE	REL	Free unused space at the end of a list of data sets.	DS	DFDSS
REFRESH	REF	Display the updated list.	DS, VOL	MVS/DFP
RESHOW	RESH	Redisplay all list entries removed by the HIDE line operator.	DS, VOL, DC, SC, MC, SG	MVS/DFP
RESTORE	REST	Restore a list of data sets that have been dumped by DFDSS.	DS	DFDSS
RETURN	RETURN	Return to the primary option menu.	All	MVS/DFP
RIGHT	RI	Scroll right the specified number of columns.	DS, VOL, DC, SC, MC, SG	MVS/DFP
SAVE	SA	Save a copy of the current list in the ISPF output table.	DS, VOL, DC, SC, MC, SG	MVS/DFP
SORT	SO	Organize lists based on entries in specific data columns.	DS, VOL, DC, SC, MC, SG	MVS/DFP
TSO Commands and CLISTs		Invoke TSO commands and CLISTs.	DS, VOL, DC, SC, MC,	MVS/DFP
UP	U	Scroll backward the specified number of entries.	DS, VOL, DC, SC, MC, SG	MVS/DFP
VALIDATE	V	Check the completeness and consistency of an entire SCDS.	CDS	MVS/DFP

Figure 78 (Page 1 of 2). ISMF Line Operators

Line Operator	Minimum Abbreviation	Description	Application	Source
ALTER	AL	Change the name of a management class or storage class associated with a data set or change the attributes of an SMS class or store group.	DS, DC, SC, MC, SG,	MVS/DFP
ANALYZE	ANAL	Examine the device and/or data set on a volume to determine if any errors exist.	VOL	MVS/DFP
BROWSE	B	View a sequential data set or a member of a PDS.	DS	MVS/DFP
CLIST	CLI	Call a TSO CLIST.	DS, VOL	MVS/DFP
COMPRESS	COM	Reclaim embedded unused space from a PDS.	DS, VOL	DFDSS
CONDENSE	CON	Free unused space at the end of a data set; compress a PDS.	DS	DFHSM
CONVERTV	CONV	Convert DASD volumes into SMS or out of SMS.	VOL	DFDSS
COPY	COP	Copy one SMS class or store group to another SMS class or store group in the same or in a different SCDS; or copy a data set and volume to a DASD volume.	DS, VOL, DC, SC, MC, SG,	MVS/DFP or DFDSS
DEFRAG	DEF	Reduce free-space fragmentation on a volume.	VOL	DFDSS
DELETE	DEL	Delete an SMS class or store group, or delete an online, backup, or DFHSM-migrated data set.	DS, DC, SC, MC, SG,	MVS/DFP or DFHSM
DISPLAY	DI	Display an SMS class and its attributes.	DC, SC, MC,	MVS/DFP
DUMP	DU	Dump a data set or volume to tape, DASD, or mass store volumes.	DS, VOL	DFDSS
EDIT	E	Edit a sequential data set or member of a PDS.	DS	MVS/DFP
HALTERDS	HA	Change the number of backup versions of a data set; change frequency of backup.	DS	DFHSM
HBACKDS	HB	Create a backup version of a data set.	DS	DFHSM
HBDELETE	HBD	Delete backup versions of a data set.	DS	DFHSM
HDELETE	HDE	Delete a migrated data set.	DS	DFHSM
HIDE	HI	Remove a list entry from display.	DS, VOL, DC, MC, SC, SG,	MVS/DFP
HMIGRATE	HM	Migrate a data set to a DFHSM level one or level two volume.	DS	DFHSM
HRECALL	HRECA	Recall a data set that has been migrated by DFHSM.	DS	DFHSM

Figure 78 (Page 2 of 2). ISMF Line Operators

Line Operator	Minimum Abbreviation	Description	Application	Source
HRECOVER	HRECO	Recover a backup version of a data set.	DS	DFHSM
INIT	INI	Initialize a volume.	VOL	MVS/DFP
INSPECT	INS	Detect defects in volume track surface.	VOL	MVS/DFP
LISTSYS	LISTS	List the systems associated with a given store group.	SG	MVS/DFP
LISTVOL	LISTV	List the volumes associated with a given store group.	SG	MVS/DFP
MESSAGE	MES	Display message text for the last operation performed on a list entry.	DS, VOL, MC, SC, DC, SG	MVS/DFP
RAUTH	RAUTH	Provide remote access codes.	VOL	MVS/DFP
RELEASE	REL	Free unused space at the end of data sets.	DS, VOL	DFDSS
REPEAT	=	Repeat the last operator command that was executed.	DS, VOL, MC, DC, SC, SG	DFDSS MVS/DFP
REFORMAT	REF	Change the volume serial number or owner ID of a volume.	VOL	MVS/DFP
RESTORE	REST	Restore data sets that have been dumped by DFDSS.	DS, VOL	DFDSS
SECURITY	SE	Invoke a RACF panel to protect data sets, store classes, or management classes.	DS, SC, MC	RACF
SETCACHE	SETC	Manage storage control characteristics.	VOL	MVS/DFP
SORTREC	SOR	Invoke DFSORT to sort data set records.	DS	DFSORT
TSO Commands and CLISTs		Invoke TSO commands and CLISTs.	DS, VOL, MC, SC, DC	MVS/DFP

Appendix B. Starter Set

The Starter Set is shipped on the basic Machine Readable Material (MRM) tape and includes the SYS1.SAMPLIB. The Starter Set consists of two parts: the JCL to import the SCDS onto your system, and source for the ACS routines:

DFPSSIMP JCL to IMPORT the SCDS (includes the exported SCDS as input data)

DFPSSSGR Storage group ACS routine

DFPSSMCR Management class ACS routine

DFPSSSCR Storage class ACS routine

DFPSSDCR Data class ACS routine

MVS/ESA System Generation and Program Directory explains how to install the Starter Set on your system. After you change the system names in the base configuration and your volume serial numbers in the storage group definitions, the Starter Set is fully functional.

Objectives

The Starter Set has the following objectives:

- Illustrate the use of the SMS classes and storage groups

- Illustrate the use of the ACS routines.

Starter Set Configuration Assumptions

The Starter Set assumes that you want to make full use of system-managed storage capabilities. Therefore, all new DASD data sets are system managed unless you specify otherwise.

In the ACS routines, it is assumed that the group name for systems program-mers is 'SYS', that the group name for system administrators is 'SYSADMIN', and that the group name for data base administrators is 'DBA'.

You must back up and restore data bases via the data base backup and recovery routines. Therefore, data base data sets are separated into their own storage group, which does not permit automatic data set or volume backups. If you always perform backups when the data bases are inactive, then these backups can be enabled.

Customizing the Starter Set to Your MVS System

Consider the following items when customizing the Starter Set:

- You need to change the volume serial numbers in the the Starter Set. They must match your installation volume serial numbers.

- If you are going to have a pool containing large data sets, evaluate the size of *large* data sets in the storage group selection routine based on migration thresholds.

- Customize the filtering techniques in the storage class and management class ACS routines according to your installation naming conventions.

- Customize the filtering techniques in the storage class ACS routine to decide which system data sets can be SMS managed. Consider excluding the following:

  - Data sets that have naming requirements which require a duplicately named data set on each system. Each permanent SMS-managed data set within a complex must have a unique name and thus you could not have SYS1.PARMLIB for each system. However, if you did manage to create the multiple duplicately named data sets, you could only reference them from one system each and thus you could not fix a problem with a data set on one system from another system.

  - Data sets that are not named SYS1.* but must be cataloged in the master catalog of more than one system.

Base Configuration Information

```
SCDS NAME:  name used in IMPORT JCL
DEFAULT MANAGEMENT CLASS:  NOACTION
DEFAULT UNIT:  3380
DEFAULT DEVICE GEOMETRY:
 • BYTES/TRACK:  47476
 • TRACKS/CYLINDER:  15
SYSTEM NAMES:  SYSTEM01
```

Starter Set Services and ACS Routines

Storage Groups

STORAGE GROUP DEFAULTS

- AUTO MIGRATE: YES
- AUTO BACKUP: YES
- AUTO DUMP: NO
- THRESHOLDS: (Non-blank)
- DUMP CLASS: (Blank)

PRIMARY   This group of volumes is where most data sets will be stored.

Note: It is necessary to change the names of the volumes used in this starter set—SMS001, SMS002, SMS003, SMS004—to correspond to the names that the customer installation actually plans to use.

- TYPE: POOL
- VOLUME LIST: SMS001, SMS002, SMS003, SMS004
- SYSTEM LIST: SYSTEM01
- THRESHOLD: (90,80)
- AUTO DUMP: YES
- DUMP CLASS: SUNDAY

LARGE   This special group of volumes is reserved for large data sets. A large amount of the volume space is made available for the allocation of large data sets.

Note: It is necessary to change the names of the volumes used in this starter set—SMS401, SMS402, SMS403, SMS404—to correspond to the names that the customer installation actually plans to use.

- TYPE: POOL
- VOLUME LIST: SMS401, SMS402, SMS403, SMS404
- SYSTEM LIST: SYSTEM01
- AUTOMIGRATION: YES
- THRESHOLD: (50,40)
- AUTO-BACKUP: YES
- AUTO-DUMP: YES
- DUMP CLASS: SUNDAY

DATABASE   This group of volumes provides isolation for the data base data sets.

Note: It is necessary to change the names of the volumes used in this starter set—SMS501, SMS502, SMS503, SMS504—to correspond to the names that the customer installation actually plans to use.

- TYPE: POOL
- VOLUME LIST: SMS501, SMS502, SMS503, SMS504
- SYSTEM LIST: SYSTEM01
- AUTOMIGRATION: YES
- THRESHOLD: (90,80)
- AUTO-BACKUP: NO
- AUTO-DUMP: NO

VIO   This is the VIO Storage Group.

Note: It is very likely that the customer installation might want to tailor the VIOMAXSIZE to meet any special needs of their installation.

- TYPE: VIO
- VIOMAXSIZE: 1024
- VIOUNITNAME: 3380
- SYSTEM LIST: SYSTEM01

Management Classes

MANAGEMENT CLASS DEFAULTS

- EXPIRE AFTER DAYS NON-USAGE: No limit
- EXPIRE AFTER DATE/DAYS: No limit
- RETENTION LIMIT: No limit
- PARTIAL RELEASE: No
- PRIMARY DAYS NON-USAGE: 2
- LEVEL 1 DAYS NON-USAGE: 60
- COMMAND OR AUTO MIGRATE: Both
- # GDG ELEMENTS ON PRIMARY: (blank)
- ROLLED-OFF GDG ACTION: (blank)
- BACKUP FREQUENCY: 1
- # OF BACKUP VERSIONS: 2

- # OF BACKUP VERSIONS: 1 (DS deleted)
- RETAIN DAYS ONLY BACKUP: 60
- RETAIN DAYS EXTRA BACKUPS: 30
- ADMIN OR USER COMMAND BACKUP: Both
- AUTO BACKUP: Yes

INTERIM   This Management Class provides for short term storage that expires and deletes data sets. Short lived GDG's should also be assigned to INTERIM. This management class is useful for storing listings that don't have any value after a few days.

- DESCRIPTION: Short life data sets
- EXPIRE AFTER DATE/DAYS NON-USAGE: 7
- # GDG ELEMENTS ON PRIMARY: 1
- ROLLED-OFF GDS ACTION: EXPIRE
- PARTIAL RELEASE: YES
- AUTO BACKUP: NO

DBSTNDRD   This Management Class provides a no-migrate and no-backup environment for data base products. It basically tells DFHSM not to manage these data sets.

- DESCRIPTION: Non-migratable data bases
- COMMAND OR AUTO MIGRATE: NONE
- AUTO BACKUP: NO

DBMIG   This Management Class provides migration, but no backup for data base products that can tolerate migration. DB2 is an example of such a data base product.

- DESCRIPTION: Migratable data bases
- PRIMARY DAYS NON-USAGE: 15
- AUTO BACKUP: NO

STANDARD   This Management Class will be used for most data sets.

- DESCRIPTION: Standard processing
- PARTIAL RELEASE: YES
- PRIMARY DAYS NON-USAGE: 15
- BACK-UP FREQUENCY: 0
- AUTO BACKUP: YES

GDG   This Management Class is for Generation Data Groups.

- DESCRIPTION: Generation Data Groups
- PARTIAL RELEASE: YES
- PRIMARY DAYS NON-USAGE: 15
- LEVEL 1 DAYS NON-USAGE: 30
- # GDG ELEMENTS ON PRIMARY: 1
- ROLLED-OFF GDS ACTION: EXPIRE
- AUTO BACKUP: NO

MONTHMIG   This Management Class provides for later than STANDARD migration delay (35 days vs 15 days).

- DESCRIPTION: Late migration data sets
- PARTIAL RELEASE: YES
- PRIMARY DAYS NON-USAGE: 35

NOMIG   This Management Class does not allow migration.

- DESCRIPTION: Non-migratable data sets
- COMMAND OR AUTO MIGRATE: NONE
- BACKUP FREQUENCY: 0

EXTRABAK   This Management Class provides for extra (5 in total) backup copies.

- DESCRIPTION: Extra backup versions data sets
- PARTIAL RELEASE: YES
- PRIMARY DAYS NON-USAGE: 15
- BACKUP FREQUENCY: 0
- BACKUP VERSIONS: 5

NOACTION   NOACTION is included as a fail-safe management class in the event that an error occurs in the coding of the management class ACS routine.

- DESCRIPTION: No processing
- COMMAND OR AUTO MIGRATE: NONE
- ADMIN OR USER COMMAND BACKUP: NONE

Storage Classes

STORAGE CLASS DEFAULTS

- AVAILABILITY: STANDARD

- GUARANTEED SPACE: NO

DBLOG   This includes any kind of logging or audit trail data set. Examples include the IMS log and the CICS journal.

- DESCRIPTION: Data base log data sets
- Direct Access
- Bias: READ
- Response: 5 ms
- Sequential
- Bias: WRITE
- Response: 3 ms
- Guaranteed Space: YES

DBBASE   This is the standard response Storage Class for data bases.

- DESCRIPTION: Standard performance data bases
- Guaranteed Space: YES

DBENHANC   This data base Storage Class provides faster response time than does DBBASE.

- DESCRIPTION: Enhanced performance data bases
- Direct Access
- Bias: READ
- Response: 10 ms
- Sequential
- Bias: READ
- Response: 10 ms
- Guaranteed Space: YES

DBCRIT   This data base Storage Class provides for continuous avail-
             ability, and the same response time as DBENHANC.

             - DESCRIPTION: High availability data bases
             - Direct Access
             - Bias: READ
             - Response: 10 ms
             - Sequential
             - Bias: READ
             - Response: 10 ms
             - Guaranteed Space: YES
             - Availability: CONTINUOUS

DBFAST   This data base Storage Class provides the fastest response
             time of all of the data base storage classes.

             - DESCRIPTION: High performance data bases
             - Direct Access
             - Bias: READ
             - Response: 5 ms
             - Sequential
             - Bias: READ
             - Response: 5 ms
             - Guaranteed Space: YES

BASE   This is the Storage Class that will be used most of the time.

             - DESCRIPTION: Standard performance data sets

ENHANCED   This Storage Class provides faster response time than does
             BASE.

             - DESCRIPTION: Enhanced performance data sets
             - Direct Access
             - Response: 10 ms
             - Sequential
             - Response: 10 ms

CRITICAL   This Storage Class provides for CONTINUOUS availability and
             the same response time as ENHANCED.

             - DESCRIPTION: High availability data sets
             - Direct Access
             - Response: 10 ms
             - Sequential
             - Response: 10 ms
             - Availability: CONTINUOUS

FASTREAD   This Storage Class provides for the fastest read response time.

             - DESCRIPTION: High read performance data sets
             - Direct Access
             - Bias: READ
             - Response: 5 ms
             - Sequential
             - Bias: READ
             - Response: 5 ms

FASTWRIT   This Storage Class provides for the fastest write response time.

- DESCRIPTION: High write performance data sets
- Direct Access
- Bias: WRITE
- Response: 5 ms
- Sequential
- Bias: WRITE
- Response: 5 ms

Data Classes

DATA CLASS DEFAULTS

- VOLUME COUNT = 1

DATAF   This is for card image type data. You can also use it for programs in source, although SRCFLIB is a more convenient way to store program source.

- DESCRIPTION: Fixed length record, sequential data sets
- RECFM: FB
- LRECL: 80
- AVGREC: U
- AVG VALUE: 80
- SPACE PRIMARY: 5000
- SECONDARY: 5000

DATAV   This is for data with a variable length record format. You can also use it for SCRIPT data sets, CLISTs, and program source. SRCVLIB is a better choice for these last three items because of the added ease of use that ISPF provides with partitioned data sets.

- DESCRIPTION: Variable length record, sequential data sets
- RECFM: VB
- LRECL: 255
- AVGREC: U
- AVG VALUE: 255
- SPACE PRIMARY: 5000
- SECONDARY: 5000

SRCFLIB   This is for storing program source code where the compiler expects the source to be of a fixed record length. Examples include FORTRAN and COBOL.

- DESCRIPTION: Fixed length record, partitioned data sets
- RECFM: FB
- LRECL: 80
- AVGREC: U
- AVG VALUE: 80
- SPACE PRIMARY: 5000
- SECONDARY: 5000
- DIRECTORY: 62

SRCVLIB   This is for storing program source code where the compiler allows the source to be of a varying record length. An example is PL/I. SCRIPT data sets and CLISTs can also use this DATACLAS.

- DESCRIPTION: Variable length record, partitioned data sets
- RECFM: VB
- LRECL: 255
- AVGREC: U
- AVG VALUE: 255
- SPACE PRIMARY: 5000
- SECONDARY: 5000
- DIRECTORY: 62

LOADLIB   This is for storing load modules. The space value, 23476, illustrates the allocation of half-tracks of 3380 via a specification in bytes. The value 23476 does not represent the block size, which is selected by the linkage editor.

- DESCRIPTION: Partitioned, load module data sets
- RECFM: U
- AVGREC: U
- AVG VALUE: 23476
- SPACE PRIMARY: 50
- SECONDARY: 50
- DIRECTORY: 62

LISTING   This is for storing output listings.

- DESCRIPTION: Sequential, listing data sets
- RECFM: VBA
- LRECL: 137
- AVGREC: U
- AVG VALUE: 137
- SPACE PRIMARY: 20000
- SECONDARY: 20000

ENTRY   A VSAM Entry Sequence Data Set. Items such as CISIZE, KEYLENGTH, and key offset are expected to be on JCL or the TSO/E command.

- DESCRIPTION: VSAM entry sequenced data sets
- RECORG: ES
- AVGREC: U
- AVG VALUE: 4096
- SPACE PRIMARY: 100
- SECONDARY: 100

KEYED   A VSAM Keyed Sequence Data Set.

- DESCRIPTION: VSAM key sequenced data sets
- RECORG: KS
- KEYOFF: 0
- AVGREC: U
- AVG VALUE: 4096
- SPACE PRIMARY: 100
- SECONDARY: 100

DIRECT   A VSAM Relative Record Data Set which allow direct access to any of its records by relative record number.

- DESCRIPTION: VSAM relative record data sets
- RECORG: RR
- AVGREC: U
- AVG VALUE: 4096
- SPACE PRIMARY: 100
- SECONDARY: 100

LINEAR   A VSAM Linear Data Set.

- DESCRIPTION: VSAM linear data sets
- RECORG: LS
- AVGREC: U
- AVG VALUE: 4096
- SPACE PRIMARY: 100
- SECONDARY: 100

ACS Routines

Storage Group

This sample storage group ACS routine provides the following:

- If the storage class indicates that the data set is a data base data set, assign it to the DATABASE storage group to prevent automatic backup of data base data sets by DFHSM.

- Assign very large data sets, determined by individual installation standards, to the LARGE storage group.

- All other data sets are assigned to the PRIMARY group.

- Try to assign all temporary data sets to the VIO storage group if they fit. Otherwise they will go to the PRIMARY storage group.

```
PROC 1 STORGRP

/*********************************************************************/
/* ACS ROUTINE TO SELECT STORAGE GROUP                               */
/*********************************************************************/

SELECT

/*********************************************************************/
/* ASSIGN DATA BASES BASED ON STORAGE CLASS                          */
/*********************************************************************/

  WHEN (&STORCLAS = DB*)
    SET &STORGRP = 'DATABASE'

/*********************************************************************/
/* ASSIGN LARGE DATA SETS BASED ON MAXIMUM SIZE                      */
/*********************************************************************/
/* NOTE:  THIS MUST BE CUSTOMIZED TO THE CORRECT VALUE               */
/* BY THE CUSTOMER SYSTEM ADMINISTRATOR.                             */
/*********************************************************************/

  WHEN (&MAXSIZE > 9999999KB)
    SET &STORGRP = 'LARGE'

/*********************************************************************/
/* ASSIGN TEMPORARY DATA SETS TO VIO AND PRIMARY STORAGE GROUPS      */
/* (SMS WILL DETERMINE VIO CANDIDACY BASED ON VIOMAXSIZE IN          */
/* VIO STORAGE GROUP)                                                */
/*********************************************************************/

  WHEN (&DSTYPE = 'TEMP')
    SET &STORGRP = 'VIO','PRIMARY'

/*********************************************************************/
/* ASSIGN EVERY OTHER DATA SET TO THE PRIMARY STORAGE GROUP          */
/*********************************************************************/

  OTHERWISE
    SET &STORGRP = 'PRIMARY'

END
END
```

Figure 79. Storage Group Routine

Management Class

This sample management class ACS routine provides the following:

- When a data set is Recalled or Recovered by DFHSM, do not change the management class if one already exists.

- Allow system administrators and systems programmers to specify a management class during allocation.

- Allow data base administrators to specify data base management classes during allocation.

- All data sets of the form SYSn.*...* allocated by the system administrator or systems programmer are assigned the NOMIG management class. These are system data sets that should not be migrated.

- Database 2 table space data sets are assigned to the DBMIG management class to allow migration of these data sets.

- Listing data sets, determined by the low level qualifier of the data set name, are assigned to the INTERIM management class.

- GDGs, determined by the data set type variable, are assigned to the GDG management class.

- Source data sets, determined by the low level qualifier of the data set name, are assigned to the EXTRABAK management class.

- Assign all other data sets to the STANDARD management class.

```
PROC 1 MGMTCLAS
/****************************************************************/
/* ACS ROUTINE TO SELECT MANAGEMENT CLASS                      */
/****************************************************************/

/****************************************************************/
/* DEFINE SHORT LIFE DATA SET LOW LEVEL QUALIFIER              */
/****************************************************************/

FILTLIST LISTING INCLUDE('SYSOUT',LIST*,'OUTLIST')

/****************************************************************/
/* DEFINE NO MIGRATION DATA SET HIGH LEVEL QUALIFIER           */
/****************************************************************/

FILTLIST NOMIG    INCLUDE(SYS%.**)

/****************************************************************/
/* DEFINE DATABASE 2 TABLE SPACE DATA SET NAME FILTER          */
/****************************************************************/

FILTLIST DB2      INCLUDE(*.DSNDB*.*.*.I0001.A*)

/****************************************************************/
/* DEFINE SOURCE DATA SETS FOR EXTRA BACKUP                    */
/****************************************************************/

FILTLIST SOURCE  INCLUDE('ASM',COB*,FOR*,'CNTL','JCL','PLI','PL1',
   'SCRIPT','CLIST','C')
```

Figure 80 (Part 1 of 2). Management Class Routine

```
SELECT

/*******************************************************************/
/* EXCLUDE DATA SET RECALL AND RECOVER FROM PROCESSING IF          */
/* MANAGEMENT CLASS IS ALREADY ASSIGNED                            */
/*******************************************************************/

  WHEN ((&ACSENVIR='RECALL' | &ACSENVIR='RECOVER') && &MGMTCLAS ¬='')
    SET &MGMTCLAS = &MGMTCLAS

/*******************************************************************/
/* ALLOW STORAGE ADMINISTRATORS AND SYSTEM PROGRAMMERS TO          */
/* SELECT THE MANAGEMENT CLASS OF CHOICE                           */
/*******************************************************************/

  WHEN ((&GROUP = 'SYSADMIN' | &GROUP = 'SYS') && MGMTCLAS ¬= ' ')
    SET &MGMTCLAS = &MGMTCLAS

/*******************************************************************/
/* ALLOW DATA BASE ADMINISTRATORS TO SELECT A DATA BASE            */
/* MANAGEMENT CLASS OF CHOICE                                      */
/*******************************************************************/

  WHEN (&GROUP = 'DBA' && &MGMTCLAS = DB*)
    SET &MGMTCLAS = &MGMTCLAS

/*******************************************************************/
/* ASSIGN NON MIGRATABLE MANAGEMENT CLASS BASED ON NAME            */
/*******************************************************************/

  WHEN (&DSN = &NOMIG)
    SET &MGMTCLAS='NOMIG'

/*******************************************************************/
/* ASSIGN DATA BASE 2 DATA SETS BY NAME                            */
/*******************************************************************/

  WHEN &DSN = &DB2
    SET &MGMTCLAS = 'DBMIG'

/*******************************************************************/
/* ASSIGN SHORT LIFE DATA SETS BASED ON LOW LEVEL QUALIFIER        */
/*******************************************************************/

  WHEN (&LLQ = &LISTING)
    SET &MGMTCLAS='INTERIM'

/*******************************************************************/
/* ASSIGN GENERATION DATA GROUPS                                   */
/*******************************************************************/

  WHEN (&DSTYPE = 'GDG')
    SET &MGMTCLAS='GDG'

/*******************************************************************/
/* ASSIGN SOURCE DATA SETS BASED ON LOW LEVEL QUALIFIER            */
/*******************************************************************/

  WHEN (&LLQ = &SOURCE)
    SET &MGMTCLAS='EXTRABAK'

/*******************************************************************/
/* ASSIGN ALL OTHER DATA SETS TO THE STANDARD MANAGEMENT CLASS     */
/*******************************************************************/

  OTHERWISE
    SET &MGMTCLAS='STANDARD'
  END
END
```

Figure 80 (Part 2 of 2). Management Class Routine

Storage Class

This sample storage class ACS routine provides the following:

- Exclude allocations to non-DASD or nontape units.

- Exclude system data sets. These are not to be SMS managed.

- Tape allocations without DATACLAS. For tape allocations, explicitly speci-
fying a data class means the data set should now be allocated as a DASD
data set.

- When a data set is recalled or recovered by DFHSM, do not change the
storage class if one already exists.

- Allocations performed by either system administrators or systems program-
mers can be excluded from system-managed storage by coding a storage
class of NONSMS.

- Allow system administrators and system programmers to specify a storage
class on their data set allocations.

- Allow data base administrators to specify storage classes beginning with
"DB."

- If the low level qualifier indicates that this data set is a data base data set,
or if it is a Database 2 table space data set, and it is being allocated by the
data base administrator, assign it a storage class of DBBASE.

- All other data sets are assigned a storage class of BASE.

```
PROC 1 STORCLAS                                                           0010021
                                                                         0020021
/*******************************************************************/    0030021
/* ACS ROUTINE TO SELECT STORAGE CLASS                          */       0040021
/*******************************************************************/    0050021
                                                                         0060021
/*******************************************************************/    0070021
/* DEFINE DASD UNIT NAMES                                       */       0080028
/*******************************************************************/    0090021
FILTLIST DASD_UNITS INCLUDE ('2305','3330','3340','3350','3375',         0110028
              '3380','SYSDA','SYSALLDA','')                              0120028
/*******************************************************************/    0220021
/* DEFINE TAPE UNIT NAMES                                       */       0230021
/*******************************************************************/    0240021
                                                                         0250021
FILTLIST TAPE_UNITS INCLUDE(TAPE*,T3420*,T3480*,'3420','3480')           0260028
                                                                         0270021
/*******************************************************************/    0280021
/* DEFINE NON SMS SYSTEM DATASETS                               */       0290021
/* ONLY SYS1 DATA SETS IN THE EXCLUDE LIST WILL BE ELIGIBLE TO BE */     0291028
/* SMS MANAGED                                                  */       0292028
/*******************************************************************/    0300021
                                                                         0310021
FILTLIST SYSDSET INCLUDE (SYS%.*)                                        0320021
                 EXCLUDE (SYS%.LIST.THAT.CAN.BE.SMS.MANAGED)             0330028
                                                                         0331028
/*******************************************************************/    0340021
/* DEFINE DATABASE 2 TABLE SPACE DATA SET NAME FILTER           */       0350021
/*******************************************************************/    0360021
                                                                         0370021
FILTLIST DB2  INCLUDE(*.DSNDB*.*.*.I0001.A*)                             0380021
                                                                         0390021
/*******************************************************************/    0400021
/* DEFINE DATA BASE LOW LEVEL QUALIFIER NAMES                   */       0410021
/*******************************************************************/    0420021
                                                                         0422029
FILTLIST DBDSET INCLUDE(DB*,'DATABASE',*DB)                              0423029
                                                                         0430021
/*******************************************************************/    0431029
/* DEFINE HSM DATA SET NAMES                                    */       0432029
/*******************************************************************/    0433029
                                                                         0434029
FILTLIST &HSMLQ1 INCLUDE('DFHSM','HSM')                                  0435029
FILTLIST &HSMLQ2 INCLUDE('SMALLDS','HMIG','BACK','VCAT','VTOC',
'DUMPVTOC')                                                              0436029
                                                                         0437029
```

Figure 81 (Part 1 of 3). Storage Class Routine

```
SELECT                                                              0530021
                                                                    0540021
/*******************************************************************/ 0620021
/* EXCLUDE NON-DASD AND NON-TAPE ALLOCATIONS                     */ 0630021
/*******************************************************************/ 0640021
                                                                    0650021
     WHEN (&UNIT ¬= &DASD_UNITS && &UNIT ¬= &TAPE_UNITS)            0660021
        SET &STORCLAS = ''                                          0670021
                                                                    0680021
/*******************************************************************/ 0690021
/* EXCLUDE SYSTEM DATA SETS THAT SHOULD NOT BE SMS MANAGED       */ 0700028
/*******************************************************************/ 0710021
                                                                    0720021
     WHEN (&DSN = &SYSDSET)                                         0730021
        SET &STORCLAS = ''                                          0740021
                                                                    0750021
/*******************************************************************/ 0760021
/* EXCLUDE TAPE ALLOCATIONS WITHOUT A DATA CLASS                 */ 0770021
/*******************************************************************/ 0780021
                                                                    0790021
     WHEN (&UNIT = &TAPE_UNITS && &DATACLAS = '')                   0800021
        SET &STORCLAS = ''                                          0810021
                                                                    0820021
/*******************************************************************/ 0830021
/* EXCLUDE DFHSM RECALLS AND RECOVERIES FROM PROCESSING IF       */ 0840021
/* STORAGE CLASS ALREADY EXISTS FOR THE DATA SET.               */ 0850021
/*******************************************************************/ 0860021
                                                                    0870021
     WHEN ((&ACSENVIR='RECALL' | &ACSENVIR='RECOVER') && &STORCLAS ¬='')  0880021
        SET &STORCLAS = &STORCLAS                                   0890028
                                                                    0900021
/*******************************************************************/ 0910021
/* ALLOW SYSTEM ADMINISTRATORS AND PROGRAMMERS TO CREATE         */ 0920021
/* NON-SMS-MANAGED DATA SETS                                     */ 0930021
/*******************************************************************/ 0940021
                                                                    0950021
     WHEN ((&GROUP = 'SYS' | &GROUP = 'SYSADMIN') && &STORCLAS='NONSMS')  0960021
        SET &STORCLAS = ''                                          0970021
                                                                    0980021
/*******************************************************************/ 0990021
/* ALLOW STORAGE ADMINISTRATORS AND SYSTEM PROGRAMMERS TO SELECT */ 1000028
/* A STORAGE CLASS OF CHOICE                                     */ 1010028
/*******************************************************************/ 1020021
                                                                    1030021
     WHEN ((&GROUP = 'SYS' | &GROUP = 'SYSADMIN') && &STORCLAS ¬= '')    1040021
        SET &STORCLAS = &STORCLAS                                   1050028
                                                                    1060021
/*******************************************************************/ 1070021
/* ALLOW DATA BASE ADMINISTRATORS TO SELECT A DATA BASE STORAGE  */ 1080021
/* CLASS OF CHOICE                                               */ 1090021
/*******************************************************************/ 1100021
                                                                    1110021
     WHEN (&GROUP = 'DBA' && &STORCLAS = DB*)                       1120021
        SET &STORCLAS = &STORCLAS                                   1130028
                                                                    1140021
/*******************************************************************/ 1150021
/* ASSIGN DATA BASE DATA SETS BASED ON DATA SET NAME             */ 1160021
/*******************************************************************/ 1170021
                                                                    1180021
     WHEN (&LLQ = &DBDSET | (&DSN = &DB2 && &GROUP = 'DBA'))   1190021
        SET &STORCLAS = 'DBBASE'                                    1200021
```

Figure 81 (Part 2 of 3). Storage Class Routine

```
                                                                                1210021
/********************************************************************************/  1220021
/* ASSIGN ALL OTHER DATA SETS TO THE BASE STORAGE CLASS                      */  1230028
/********************************************************************************/  1240021
                                                                                1250021
     OTHERWISE SET &STORCLAS = 'BASE'                                            1260021
                                                                                1270021
     END                                                                         1280021
                                                                                1280129
/********************************************************************************/  1281029
/* ITEMS IN FILTLIST FOR &HSMLQ1 ARE THE VALUES SPECIFIED ON THE            */  1282029
/* SETSYS BACKUPPREFIX OR MIGRATEPREFIX PARAMETERS, OR IF THESE ARE         */  1282129
/* NOT SPECIFIED, THE SPECIFIED (OR DEFAULT) VALUE ON THE UID               */  1282229
/* PARAMETER OF THE DFHSM STARTUP PROCEDURE.                                */  1282329
/********************************************************************************/  1283029
                                                                                1284029
     IF &DSN(1) = &HSMLQ1 AND                                                    1285029
        &DSN(2) = &HSMLQ2 THEN                                                   1286029
           SET &STORCLAS = ''                                                    1287029
                                                                                1288029
END                                                                             1290021
                                                                                1300024
```

Figure 81 (Part 3 of 3). Storage Class Routine

Data Class

This sample routine chooses Data Classes solely based on the lowest level qualifier.

```
PROC 1 DATACLAS
/********************************************************************************/
/* ACS ROUTINE TO SELECT DATA CLASS                                          */
/********************************************************************************/

/********************************************************************************/
/* DEFINE RECOGNIZED LOW LEVEL DATA SET NAME QUALIFIERS                       */
/********************************************************************************/

FILTLIST DATAF   INCLUDE('DATA','FDATA')
FILTLIST DATAV   INCLUDE('TEXT','VDATA')
FILTLIST SRCF    INCLUDE(COB*,FOR*,'JCL') EXCLUDE(**.SPFTEMP%.CNTL)
FILTLIST SRCV    INCLUDE('PLI','PL1','SCRIPT','CLIST')
FILTLIST OBJ     INCLUDE('SYSGO','SYSLIN',OBJ*)
FILTLIST LOAD    INCLUDE('RESLIB','PGMLIB',LOAD*)
FILTLIST LIST    INCLUDE('SYSOUT',LIST*,'OUTLIST')

/********************************************************************************/
/* IF THE ALLOCATION SPECIFIES A DATA CLASS, THEN ALLOW IT TO BE            */
/* USED.  OTHERWISE, SELECT A DATA CLASS BASED ON THE LOW LEVEL             */
/* QUALIFIER.                                                               */
/********************************************************************************/

IF &DATACLAS='' THEN
  SELECT (&LLQ)
     WHEN (&DATAF)     SET &DATACLAS='DATAF'
     WHEN (&DATAV)     SET &DATACLAS='DATAV'
     WHEN (&SRCF)      SET &DATACLAS='SRCFLIB'
     WHEN (&OBJ)       SET &DATACLAS='DATAF'
     WHEN (&SRCV)      SET &DATACLAS='SRCVLIB'
     WHEN (&LOAD)      SET &DATACLAS='LOADLIB'
     WHEN (&LIST)      SET &DATACLAS='LISTING'
     WHEN ('ENTRY')    SET &DATACLAS='ENTRY'
     WHEN ('KEYED')    SET &DATACLAS='KEYED'
     WHEN ('DIRECT')   SET &DATACLAS='DIRECT'
     WHEN ('LINEAR')   SET &DATACLAS='LINEAR'
  END
END
```

Figure 82. Data Class Routine

Appendix C. Choosing Restricted Line Operators and List Commands

This chapter shows you how to perform the following types of data and storage management tasks:

- Maintaining the media characteristics of a volume

- Modifying DASD storage control

- Providing access to IBM Service Representatives at remote locations.

To perform these tasks, you will enter line operators or list commands on data set and/or volume lists and then complete data entry panels. Refer to your online help panels or the *ISMF User's Guide*.

Maintaining the Media Characteristics of a Volume

You can use ISMF to examine and manage the characteristics of your data stored on DASD volumes. There are four line operators available to perform the following media maintenance tasks:

- Inspect a subset of a volume and check the surface of a track to determine if there is a defect. ISMF also provides options to help you remedy the problem.

- Examine a drive and your data to determine if errors exist. This function helps you distinguish between errors caused by drive problems and errors caused by media problems.

- Initialize a DASD volume for use in an OS/VS system

- Change the volume label of a DASD volume. You can specify a new volume serial number or a new owner ID.

Figure 83 on page 160 summarizes the ICKDSF functions that ISMF uses to perform these tasks.

Figure 83. Summary of Media Maintenance Line Operators

Task	Function	Scope	Action
Detect defects in volume track surface.	INSPECT	Volume	Inspects volume track surfaces for defects. Skips defects, assigns alternate tracks, reclaims tracks previously found defective, or prints a map of defective tracks. Performed in the background by ICKDSF.
Diagnose a problem to determine if is a drive error or a media error.	ANALYZE	Volume	Examines a drive and your data on a volume to determine if errors exist. Provides output designed to aid you in distinguishing between drive problems and media problems. Performed in the background by ICKDSF.
Initialize a volume.	INIT	Volume	Initializes a DASD volume for use in an OS/VS system. Performed in the background by ICKDSF.
Change the volume serial number or owner ID of a volume.	REFORMAT	Volume	Specifies a new volume serial number and/or owner ID for an existing volume. Performed in the background by ICKDSF.

Detecting Defects in Volumes

There are two ways to check a volume for defects. You can use the INSPECT line operator or the ANALYZE line operator. You can only use these line operators in the volume application or the storage group application. You must have ICKDSF Release 9 or later installed on your system.

The **INSPECT** line operator generates a job to inspect a subset of a volume. INSPECT can determine if there is a defect in a track surface. You can also specify that ISMF skip a defect, assign alternate tracks, reclaim tracks that were previously found defective, or print a map of the defective tracks.

You can enter INSPECT as a line operator from the Volume List Panel where you can choose between various options:

- The type of inspection.

- The range of the inspection. You can choose how much of the volume you would like to inspect from the following options:

    Inspect the entire volume
    Inspect a range of tracks
    Inspect particular tracks.

- The level of surface checking. You can choose whether or not defective tracks are skipped with the following options:

  Do not perform surface checking
  Perform skip displacement on those tracks where a data check is detected
  Perform skip displacement on all the tracks that are specified.

- The level of data preservation. You can specify what ISMF will do with your data while the tracks are inspected:

  Save the data in permanent storage
  Save the data in temporary storage
  Do not save the data.

The INSPECT Track Range Entry Panel and the INSPECT Discrete Track Entry Panel allow you to specify track ranges or particular tracks for the inspection. Once you specify the tracks that you would like to inspect, ISMF takes you to the INSPECT Options Entry Panel. On this panel, you can specify that ISMF perform any of the following functions:

- Assign alternate tracks to the defective primary tracks.

- Print a volume map that lists the tracks that are defective or inconsistent.

- Check each track a particular number of times.

- Reclaim tracks that were flagged as defective if they appear to be usable after surface checking.

- Inspect an offline volume that has been formatted as a Mass Storage Subsystem staging pack.

- Check a specified data set for the passwords of the volumes that are password protected.

- Continue the inspection or terminate the inspection of a data set that cannot be accessed exclusively.

Once you complete the INSPECT Options Entry Panel, ISMF displays the INSPECT Job Submission Entry Panel. From this panel, you can submit the job or specify a data set where ISMF will save the job. The Job Submission Entry Panel also optionally allows you to change the JCL that ISMF will use to run the job.

The **ANALYZE** line operator generates a job to examine a drive and/or your data on a volume to help determine if errors exist. The output from the ANALYZE line operator is designed to help you distinguish between drive problems and media problems. This output can also help you locate and fix these problems. ANALYZE has two main functions: a drive test and a data verification test.

The ANALYZE drive test performs fundamental tests on the device hardware. These tests ensure that the hardware can perform basic operations such as seek, read, and write. The drive test is only valid on IBM DASDs that have nonremovable storage media.

The ANALYZE data verification test reads the data records on a volume. This test does not transfer your data to perform the operation. You can ANALYZE the entire volume, a range of tracks, or specific tracks.

You can enter the ANALYZE line operator from the Volume List Panel. ISMF takes you to the ANALYZE Entry Panel. From this panel you can choose between the various options. Additional pages for the ANALYZE Entry Panel allow you to specify track ranges or specific tracks.

Once you complete the ANALYZE Entry Panel, ISMF displays the ANALYZE Job Submission Entry Panel. From this panel, you can submit the job or specify a data set where ISMF will save the job. The Job Submission Entry Panel also optionally allows you to change the JCL that ISMF will use to run the job.

From the Entry Panel for each of these line operators, the online help provides detailed descriptions of each field.

Formatting Volumes

There are two line operators available for you to use to format a volume. The INIT line operator allows you to initialize a volume. The REFORMAT line operator allows you to specify a new volume serial number and owner ID for an existing volume. These line operators can be entered from the Volume List Panel in the volume application or the storage group application. You must have ICKDSF Release 9 or later on your system to use the REFORMAT line operator.

The **INIT** line operator allows you to initialize a volume and build the Volume Table of Contents (VTOC) and VTOC index.

When you specify INIT from the storage group application, the volume corresponds to an SMS format. When you specify INIT from the volume application, the volume corresponds to a non-SMS format.

Entering the INIT line operator takes you to the INIT Entry Panel. ISMF uses the information you specify on page 1 and 2 of the INIT Entry Panel to determine where to build the VTOC and VTOC index, and how to initialize the volume. You can initialize the volume either online or offline. ISMF provides three types of initialization:

- Minimal initialization sets up the IPL bootstrap and volume label, builds the VTOC and VTOC index.

- Medial initialization validates the home address and record 0 of each track. Medial initialization also performs the functions under minimal initialization.

- Maximal initialization checks the surface of the tracks. Maximal initialization also performs the functions of medial initialization. (Maximal initialization is not supported for 3375 and 3380 devices).

Once you complete the INIT Entry Panel, ISMF displays the INIT Job Submission Entry Panel. From this panel, you can submit the job or specify a data set where ISMF will save the job. The Job Submission Entry Panel also optionally allows you to change the JCL that ISMF will use to run the job.

You can use the **REFORMAT** line operator to generate a job that changes the volume label of a DASD volume. You can change the volume serial number and/or the owner ID.

REFORMAT is entered from the Volume List Panel. ISMF takes you to the REFORMAT Entry Panel. From this panel, you can specify a new volume serial

number for the volume, a new owner ID for the volume, or both. Your data in the volume is not affected by reformatting.

Once you complete the REFORMAT Entry Panel, ISMF displays the REFORMAT Job Submission Entry Panel. From this panel, you can submit the job or specify a data set where ISMF will save the job. The Job Submission Entry Panel also optionally allows you to change the JCL that ISMF will use to run the job.

Modifying DASD Storage Control Characteristics

You can use the **SETCACHE** line operator to manage the caching of a storage control unit. Figure 84 describes the SETCACHE line operator.

Figure 84. Summary of the SETCACHE Line Operator

Task	Function	Scope	Action
Manage storage control unit character- istics	SETCACHE	Volume and storage control unit	ISMF generates a job that modifies the caching status, duplexing status, and also causes data to be destaged in the storage control units associated with specific DASD volumes. ISMF performs these functions in the foreground using a TSO command or in the background using the ISMF job sub- mission facility.

You can use the SETCACHE line operator from the volume list panel to modify the storage control unit characteristics. The following characteristics can b? modified with the SETCACHE line operator:

- Performance characteristics control the caching of read and/or write requests. You can modify these characteristics to reduce the frequency of access to DASD.

- Availability characteristics control the automatic back up of the device activity.

- Resources characteristics control how data that is stored in the cache is destaged or discarded. You can directly manage data in cache and nonvol- atile storage.

The SETCACHE line operator is entered against a particular volume from the Volume List Panel. You can invoke SETCACHE with a PERFORMANCE, AVAIL- ABILITY, or RESOURCES parameter depending on the characteristic you would like to modify. You can also enter SETCACHE without a parameter. On the IBM 3990 Storage Control Model 3, ISMF will take you to the SETCACHE Features Entry Panel. From this panel you can select the characteristic you would like to modify. ISMF will take you to either the PERFORMANCE, AVAILABILITY, or RESOURCES Entry Panel. On the IBM 3880 Storage Control Model 13/23, ISMF takes you to the PERFORMANCE Entry Panel where you can do your modifica- tions.

The specific features that you can modify depend on the model of the storage control unit that you are working with. You must also have storage adminis- tration authorization to use the SETCACHE line operator.

Modifying Caching Characteristics

The SETCACHE PERFORMANCE Entry Panel is used to modify the caching characteristics of the storage control units. If you are using an IBM 3990 Storage Control Model 3, you have support for the caching of both read and write requests. If you are using an IBM 3880 Storage Control Model 13 or 23, you have support for the caching of read requests. In the terminology of the SETCACHE line operator, the caching of read requests is controlled by the READ CACHE or READ and SYSTEM CACHE functions. The caching of write requests is controlled by the DASD FAST WRITE, NON-VOLATILE STORAGE (NVS), and CACHE FAST WRITE functions.

Read Caching/Read and System Caching

You can establish or terminate the READ CACHE or READ and SYSTEM CACHE functions to control the caching of read requests. Read caching is performed on a track basis. When you need data from a volume, a read request is made to the subsystem. The subsystem reads the data from the volume and sends it back to you. At the same time, a copy of all the data on the track that the subsystem accesses is saved in cache that resides in the subsystem. Any subsequent requests that you make for data on the track can be satisfied from the cache. Read performance is improved when the data for your read requests come from the cache.

You can specify that the storage control unit perform this type of read caching on a volume level. Caching can be turned on and off for individual volumes. This feature gives you greater flexibility to manage the volumes that are using the caching resources.

Write Caching

If your storage control unit supports the caching of write requests, you can use the SETCACHE line operator to establish or terminate the DASD FAST WRITE and CACHE FAST WRITE functions. Note that read and system caching must be active in order for fast write (both DASD and CACHE) to be in effect.

DASD FAST WRITE provides support for the caching of write requests. Nonvolatile storage (NVS) provides backup storage for this function. When you write data out to a volume, a write request is made to the subsystem. The subsystem accepts the data and puts it in the subsystem cache. Later, the storage control unit schedules the data for writing.

When the data is written to the subsystem cache, it is also written to NVS. NVS is a buffer backed up by battery power that can maintain data up to 48 hours. NVS provides insurance against the loss of data written to the subsystem cache and not yet sent to DASD.

The CACHE FAST WRITE feature allows part of the subsystem storage to be set aside as work space. Data written out in CACHE FAST WRITE mode is not staged for writing to DASD unless you make an explicit request or there is a need to re-use some of the cache space. CACHE FAST WRITE is for temporary work files only since data in cache is eventually discarded.

Modifying Duplexing Characteristics

The SETCACHE AVAILABILITY Entry Panel is used to manage the availability of a duplex pair of devices that control the duplication and recovery of data. The AVAILABILITY characteristics only affect data that is written to DASD.

To establish a duplex pair, you must specify:

- A primary and a secondary device

- The type of synchronization between the two devices.

You can specify the rate of copy in the synchronization. Once duplexing is established, the storage control unit writes data to the primary device and simultaneously places a copy of it in cache along with control information in NVS. At a later time, the data in cache is written out to the secondary device, and the control information in NVS is updated.

If duplexing is interrupted, the dual copy pair is suspended. The storage control still writes information to NVS, but data is written to only the primary device. When duplexing is resumed, the secondary device is placed back in synchronization by the storage subsystem.

Modifying Destaging Characteristics

The SETCACHE RESOURCES Entry Panel is used to directly manage the data that has already been written to the cache and NVS. This function of the SETCACHE line operator allows you to:

- Destage data

- Discard pinned data

- Set all the volumes and the subsystem back to their default status.

Once you have written data to the cache or NVS the RESOURCES Entry Panel allows you to write the data to DASD. If this destaging process fails, the data is pinned in cache or NVS. The RESOURCES Entry Panel allows you to discard this pinned data. All the volumes attached to a storage subsystem along with the subsystem itself can be set back to their default status by re-initializing the storage subsystem. Note that re-initializing a subsystem would cause all duplex pairs to be lost and any fast write data in cache and/or NVS to be discarded.

Job Submission

When you specify the SETCACHE line operator, ISMF will take you through the entry panels you need to complete given the functions you choose to perform. ISMF will eventually ask you to specify whether the job will be performed in the foreground or the background. In the foreground, ISMF uses the IDCAMS TSO SETCACHE support. ISMF will subsequently redisplay updated versions of the entry panels with their new values after the the job has executed. The list panels will not be updated unless a REFRESH is done. For some functions, execution in foreground can occupy your TSO terminal for 20 minutes or more. In these cases, execution in the background is recommended.

When you specify that the SETCACHE job be performed in background, ISMF displays the IDCAMS Job Submission Entry Panel. From this panel, you can submit the job or specify a data set where ISMF will save the job. The Job

Submission Entry Panel also optionally allows you to change the JCL that ISMF will use to run the job. At each level of the SETCACHE panels, the ISMF online help provides detailed descriptions for each of the fields.

Once ISMF takes you into the SETCACHE Entry Panel, the online help provides detail descriptions for each of the fields.

Providing Remote Authorization Codes

The RAUTH line operator provides an easy way to acquire the remote access authorization codes for your IBM 3990 Storage Control Model 3device. Figure 85 summarizes the RAUTH line operator.

Figure 85. Summary of the RAUTH Line Operator

Task	Function	Scope	Action
Provide remote access codes	RAUTH	Volume	Invokes a RAUTH display panel that shows the remote access authorization codes.

When you enter the RAUTH line operator from the Volume List Panel, ISMF returns passwords on the RAUTH Display Panel.

Each password is valid for one hour. Once the password is displayed, ISMF will not display the password again. If you specify RAUTH again, ISMF will return new passwords.

Using these passwords, an off-site IBM Service Representative can log on to all models of the IBM storage control family (Models 1, 2, and 3). The Service Representative can then provide help diagnosing and correcting problems.

You must have Storage Administrator authorization to use the RAUTH line operator. Since the standard support facilities of ISPF are available, for example, printing the screen, it is your responsibility to maintain the security of the passwords. Also, RACF DASDVOL alter authority is required over the volume the request is made against. ISMF online help provides a description of each of the fields on the RAUTH display panel.

Abbreviations

The following terms and abbreviations are defined as they are used in the MVS/DFP library. If you do not find the term or abbreviation you are looking for, see *Dictionary of Computing*, SC20-1699 (formerly published as *IBM Vocabulary for Data Processing, Telecommunications, and Office Systems*, GC20-1699).

This list includes acronyms and abbreviations developed by the American National Standards Institute (ANSI) and the International Organization for Standardization (ISO). This material is reproduced from the *American National Dictionary for Information Processing*, copyright 1977 by the Computer and Business Equipment Manufacturers American National Standards Institute, 1430 Broadway, New York, New York 10018.

ACDS. Active control data set.

ACS. Automatic class selection.

BCS. Basic catalog structure.

CDS. Control data set.

CFW. CACHE Fast Write.

DASD. Direct access storage device.

DFDSS. Data Facility Data Set Services.

DFHSM. Data Facility Hierarchical Storage Manager.

DFW. DASD Fast Write.

ISPF. Interactive System Productivity Facility.

NVS. Nonvolatile Storage

RACF. Resource Access Control Facility.

SMS. Storage Management Subsystem.

TSO. Time sharing option.

VSAM. Virtual storage access method.

VTOC. Volume table of contents.

VVDS. VSAM volume data sets.
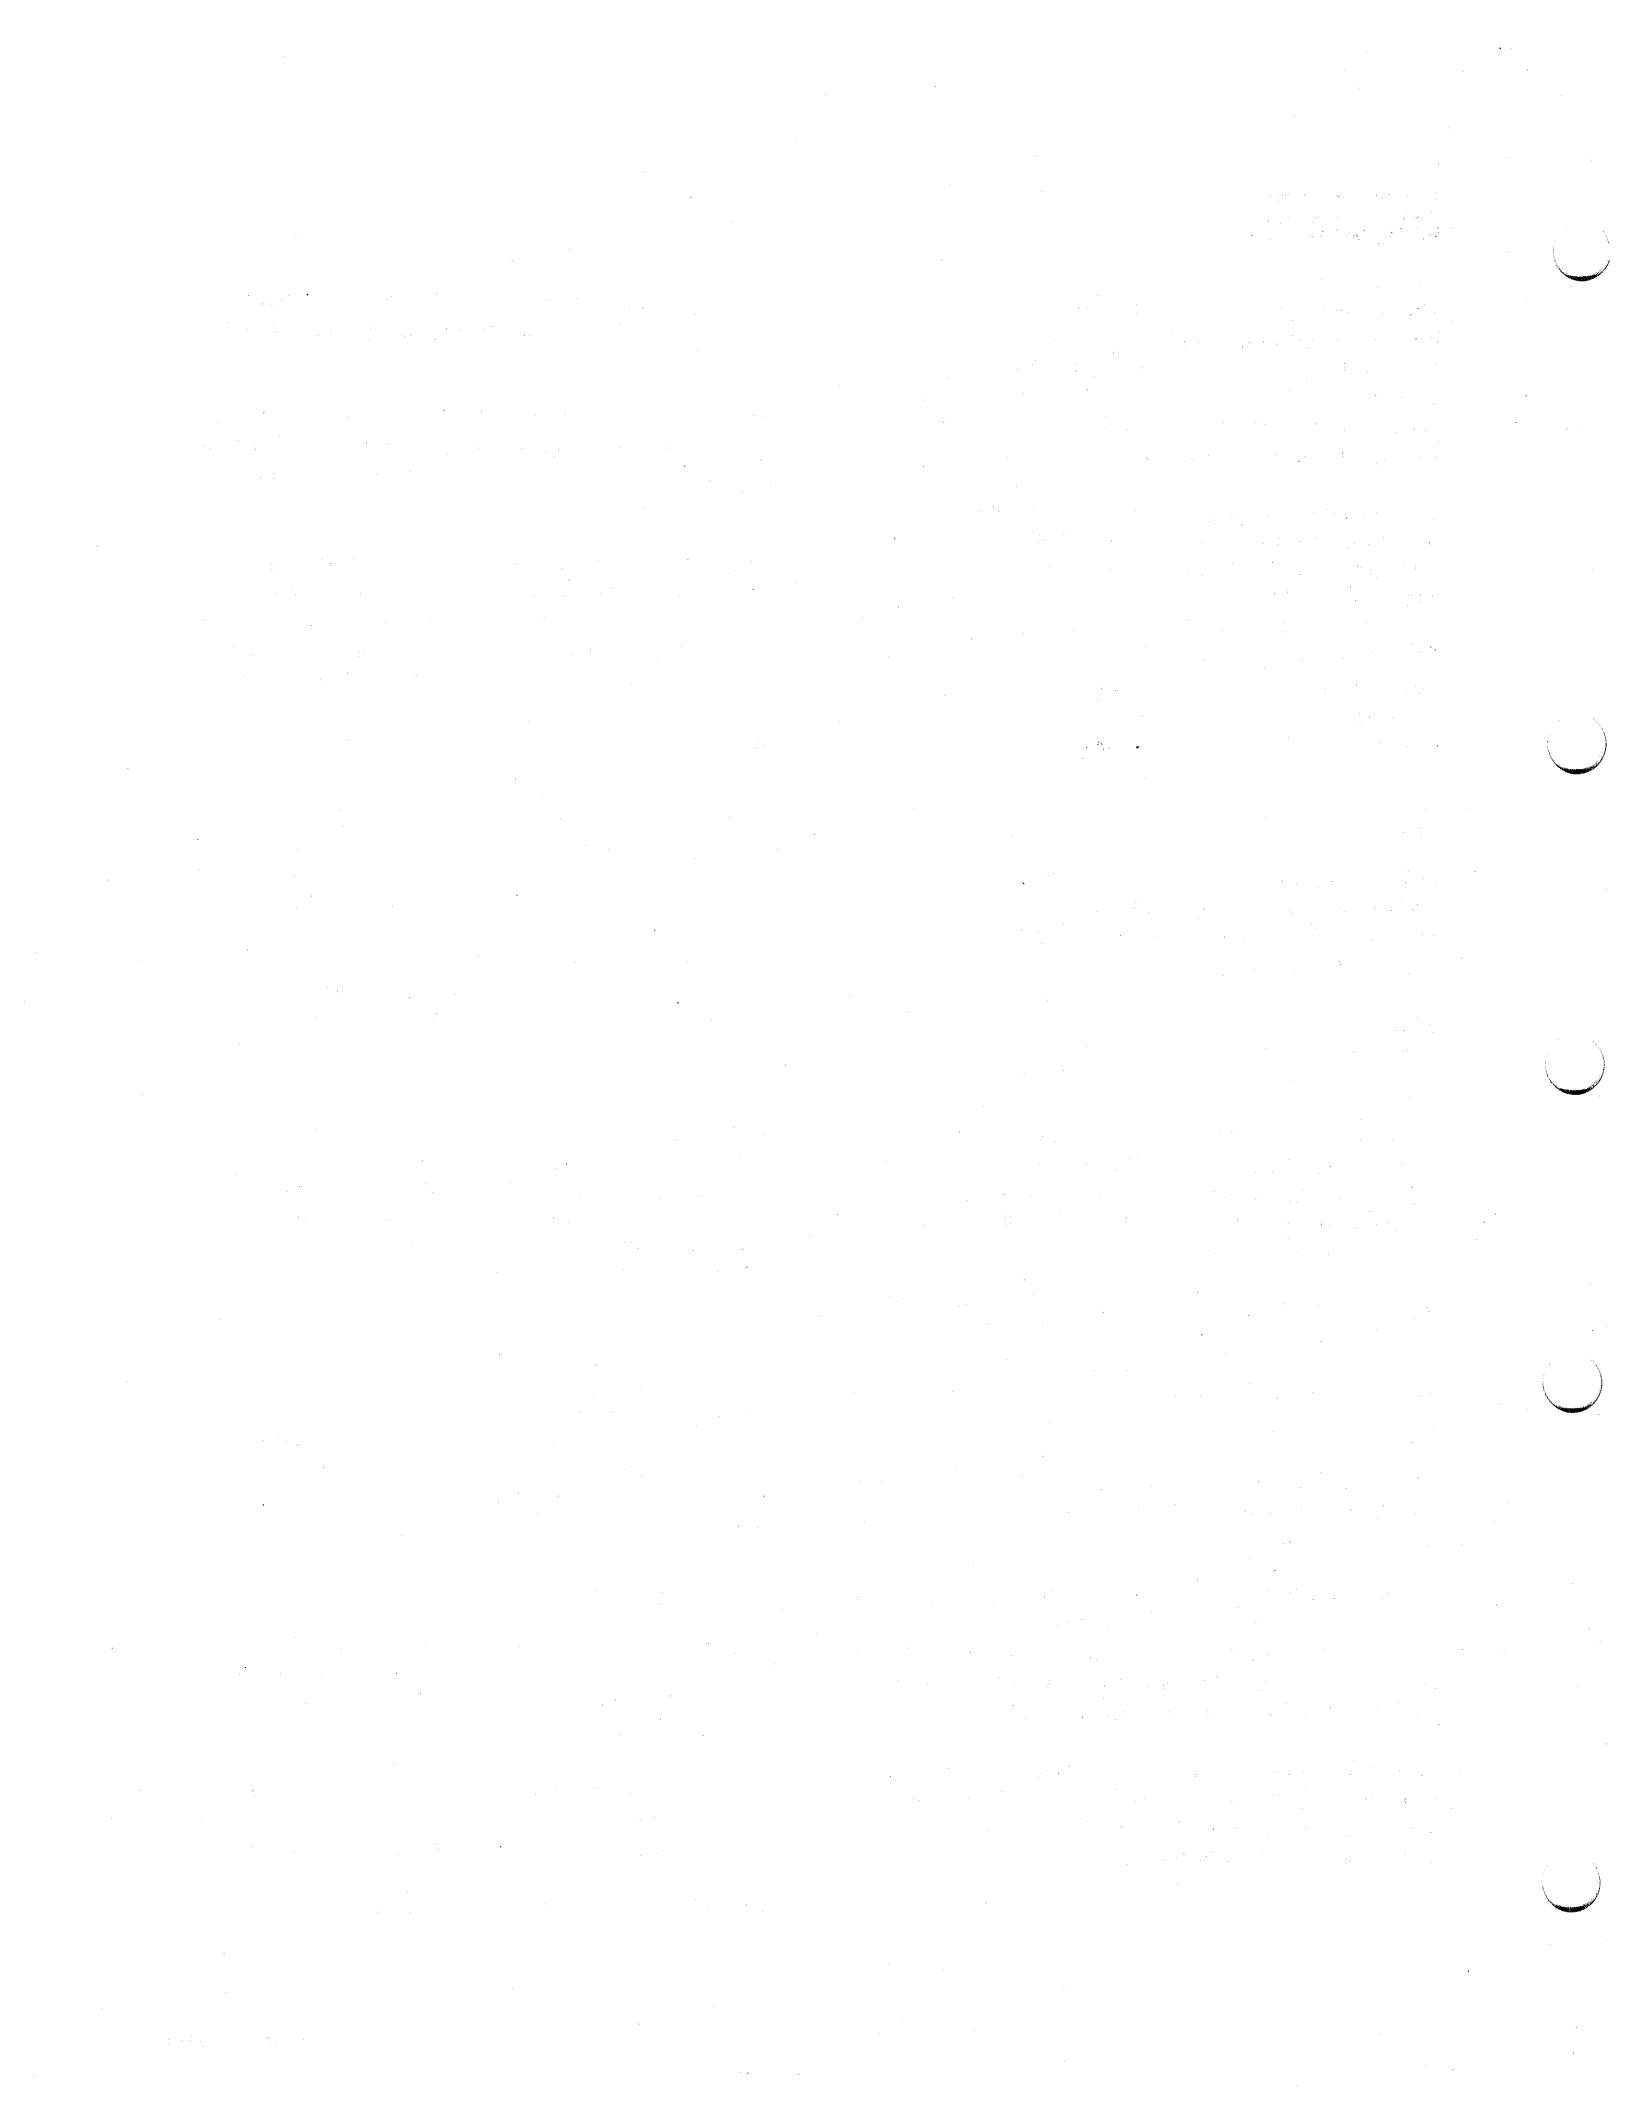
Glossary

The following terms and abbreviations are defined as they are used in the MVS/DFP library. If you do not find the term or abbreviation you are looking for, see *Dictionary of Computing*, SC20-1699 (formerly published as *IBM Vocabulary for Data Processing, Telecommunications, and Office Systems*, GC20-1699).

This glossary includes acronyms and abbreviations developed by the American National Standards Institute (ANSI) and the International Organization for Standardization (ISO). This material is reproduced from the *American National Dictionary for Information Processing*, copyright 1977 by the Computer and Business Equipment Manufacturers American National Standards Institute, 1430 Broadway, New York, New York 10018.

A

abend. Abnormal end of task: termination of a task prior to its completion because of an error condition that cannot be resolved by recovery facilities while the task is executing.

ACS installation exit. User-written code, executed after an ACS routine, that provides capabilities beyond the scope of the ACS routine.

ACS routine. A procedural set of ACS language statements. Based on a set of input variables, the ACS language statements generate the name of a predefined SMS class, or a list of names of predefined storage groups, for a data set.

activate. To load the contents of an SCDS into SMS address space storage and into an ACDS, or to load the contents of an existing ACDS into SMS address space storage. This establishes a new storage management policy for the SMS complex.

active configuration. The most recently activated SCDS, which controls storage management for the SMS complex.

active control data set (ACDS). A VSAM linear data set that contains a copy of the most recently activated configuration (SCDS) and subsequent updates. All systems in an SMS complex use the ACDS identified in IGDSMSmm to manage storage.

automatic backup. In DFHSM, the process of automatically copying eligible data sets from primary volumes or migration volumes to backup volumes during a specified backup cycle.

automatic class selection (ACS). A mechanism for assigning SMS classes and storage groups to data sets.

automatic dump. In DFHSM, the process of using DFDSS to automatically do a full volume dump of all allocated space on primary volumes to designated tape dump volumes.

automatic migration. In DFHSM, the process of automatically moving eligible data sets either from primary volumes to level 1 volumes or from level 1 volumes to level 2 volumes without a specific request for each data set so moved. See *interval migration*.

B

backup. In DFHSM, the process of copying a data set residing on a level 0 volume, level 1 volume, or a volume not managed by DFHSM to a backup volume. See *automatic backup* and *incremental backup*.

base configuration information. Part of an SMS configuration, it contains the default management class, default unit, and default device geometry. It also identifies the systems that an SMS configuration is to manage.

basic catalog structure (BCS). The name of the actual catalog structure within the integrated catalog facility environment. The integrated catalog facility catalog is composed of a BCS together with its related volume tables of contents (VTOCs) and VSAM volume data sets (VVDSs).

C

catalog. A data set that contains extensive information required to locate other data sets, to allocate and deallocate storage space, to verify the access authority of a program or operator, and to accumulate data set usage statistics.

class. See *SMS class*.

CLIST. A sequential list of commands and control statements assigned a single name; when the name is invoked the commands in the list are executed in sequential order.

cluster. A data component and an index component in a VSAM key-sequenced data set; or a data component alone in a VSAM entry-sequenced data set.

COMMDS. See *communications data set*.

communications data set (COMMDS). The primary means of communication among systems in an SMS complex. Shared among the systems in the SMS complex, the COMMDS is a VSAM linear data set that contains the name of the ACDS and current utilization statistics for each SMS-managed volume.

compaction. In DFHSM, a method of compressing and encoding data during migration or backup to reduce storage space.

complex. See *SMS complex.*

compression. A method for removing imbedded unused space from between members of a partitioned data set.

connectivity. Relationship that establishes the eligibility of a given system in an SMS complex to access a VIO storage group, a pool storage group, and the individual volumes within a pool storage group. The relationship can be NOTCON (not connected), indicating ineligibility, or any of the following, all of which imply eligibility: ENABLE, QUIALL (quiesce all), QUINEW (quiesce new), DISALL (disable all), DISNEW (disable new).

construct. A collective name for data class, storage class, management class, and storage group.

control data set (CDS). With respect to SMS, a VSAM linear data set containing configurational, operational, or communication information. SMS introduces three types of control data sets: the source control data set, the active control data set, and the communications data set.

CONVERTED. A physical volume status indicating that all of the data sets on a volume have an associated storage class and are cataloged in an integrated catalog facility catalog. SMS can select a CONVERTED volume for all supported functions. See INITIAL.

D

data class. A list of allocation attributes that the system uses for the creation of data sets.

Data Facility Data Set Services (DFDSS). An IBM licensed program used to copy, move, dump and restore data sets and volumes.

Data Facility Hierarchical Storage Manager (DFHSM). An IBM licensed program used to back up, recover, and manage space on volumes.

Data Facility Product (DFP). An IBM licensed program used to manage programs, devices, and data in an MVS operating environment.

default device geometry. Part of the configuration base, it identifies the number of bytes per track and the number of tracks per cylinder for converting TRK and CYL requests into bytes, when no generic UNIT has been specified. SMS uses the converted value to determine how many tracks or cylinders to allocate to the data set.

default management class. Part of the configuration base, it identifies the management class that DFHSM is to apply to SMS-managed data sets that do not have a management class.

default unit. Part of the configuration base, it identifies an esoteric (such as SYSDA) or generic (such as 3380) device name. If an end user fails to specify the UNIT parameter on their JCL, SMS applies the default unit to all data sets having a disposition of either MOD or NEW that are *not* SMS managed.

direct access storage device (DASD). A storage device in which the access time is effectively independent of the location of the data.

DISALL (disable all). Relationship that prevents a system from allocating or accessing data sets in a VIO storage group, a pool storage group, or individual volumes within a pool storage group.

DISNEW (disable new). Relationship that prevents a system from allocating new data sets in a VIO storage group, a pool storage group, or individual volumes within a pool storage group.

dummy storage group. A type of storage group that contains the serial numbers of volumes that no longer exist on a system but that end users continue to reference in their JCL. These volumes are treated as SMS-managed volumes. Dummy storage groups allow old JCL to function without having to be changed.

dump class. A DFHSM-named set of characteristics that describe how volume dumps are managed.

E

ENABLE. Relationship that allows a system to allocate and access data sets in a VIO storage group, a pool storage group, or individual volumes within a pool storage group.

end user. A person in a data processing installation who requires the services provided by the computer system.

esoteric name. A name used to define a group of devices having similar hardware characteristics, such as TAPE or SYSDA. See *generic name.*

extent. A continuous space on a DASD volume occupied by a data set or portion of a data set.

F

FILTLIST (filter list). An inclusion/exclusion pair of lists that define filter criteria. FILTLISTs simplify comparison operations in ACS routines.

G

generic name. A name assigned to a class of devices (such as 3380) that is derived from the IODEVICE statement in the MVS configuration program. See *esoteric name*.

guaranteed space. A storage class attribute indicating that space is to be preallocated when a data set is created. If explicit volume serial numbers are specified, SMS honors them. If space to satisfy the allocation is not available on the user-specified volumes, the allocation fails.

I

incremental backup. In DFHSM, the process of copying a data set that has been opened for other than read-only access since the last backup version was created, and that has met the backup frequency criteria.

INITIAL. A physical volume status indicating that some data sets on a given volume lack an associated storage class or are not cataloged in an integrated catalog facility catalog. An INITIAL volume is only partially converted to SMS. See CONVERTED.

in-place conversion. The process of bringing a volume and the data sets it contains under the control of SMS, without requiring a second volume to serve as an intermediary holding place during the conversion.

interactive. Pertaining to an application in which each entry calls forth a response from the system or program, usually in the form of a conversational dialog with screens or menus.

Interactive System Productivity Facility (ISPF). An IBM licensed program used to develop, test, and run application programs interactively. ISPF is the interactive interface for all storage management functions.

interval migration. In DFHSM, automatic migration that occurs when the high threshold of occupancy is reached or exceeded on a primary volume during a specified time interval. Data sets are moved from the volume, oldest first, until the low threshold of occupancy is reached.

invalid. A status indicating that an SCDS or ACDS is in error. Either the SCDS is incomplete, or an ACS routine in the SCDS has assigned an undefined SMS class name or assigned a storage group list containing undefined names. To be complete, an SCDS must contain at least one storage class definition, at least one pool storage group definition containing at least one volume, a storage group selection routine, and base configuration information. An invalid SCDS cannot be activated.

J

job control language (JCL). A problem-oriented language used to express statements in a job that identify the job or describe its requirements to an operating system.

L

level 0 volume. A primary volume or a user volume not managed by DFHSM

level 1 volume. A volume owned by DFHSM containing data sets that migrated from a level 0 volume.

level 2 volume. A volume under control of DFHSM containing data sets that migrated from a level 0 or a level 1 volume.

logical storage. With respect to data, the attributes that describe the data and its usage, as opposed to the physical location of the data.

M

management class. A list of data set migration, backup, and retention attributes that DFHSM uses to manage storage at the data set level.

master catalog. A catalog that points to user catalogs. See *catalog*.

migration. In DFHSM, the process of moving a cataloged data set from a primary volume to a level 1 volume or level 2 volume, from a level 1 volume to a level 2 volume, or from a volume not managed by DFHSM to a level 1 or level 2 volume.

MVS/DFP. An IBM licensed program which is the base for the Storage Management Subsystem.

MVS/Extended Architecture (MVS/XA). An MVS operating system environment that supports 31-bit real and virtual storage addressing, increasing the size of addressable real and virtual storage from 16 megabytes to 2 gigabytes.

N

non-SMS volume. A volume that is not controlled by SMS.

NOTCON (not connected). Relationship that indicates a system is defined but has no access to a VIO storage group, a pool storage group, or individual volumes within a pool storage group.

O

online. Pertaining to equipment, devices, or data under the direct control of the processor.

P

partitioned data set (PDS). A data set in DASD storage that is divided into partitions, called *members*, each of which can contain a program, part of a program, or data.

performance. For a storage subsystem, a measurement of effective data processing speed against the amount of resource that is consumed by a complex. Performance is largely determined by throughput, response time, and system availability.

physical storage. With respect to data, the actual space on a storage device that is to contain data.

physical volume status. Relationship between a volume and SMS. Physical volume status indicates if a volume is SMS managed, if all of its data sets have an associated storage class, and if all of its data sets are cataloged in an integrated catalog facility catalog. A volume can be in one of three states: CONVERTED (fully converted), INITIAL (partially converted), or non-SMS (unconverted).

pool. See *storage pool*

pool storage group. A type of storage group that contains the serial numbers of SMS-managed volumes. Pool storage groups allow groups of volumes to be managed as a single entity.

primary volume. A volume managed by DFHSM containing data sets that are directly accessible to the user.

Q

QUIALL (quiesce all). Relationship that prevents a system from scheduling jobs that allocate or access data sets in a VIO storage group, a pool storage group, or individual volumes within a pool storage group.

QUINEW (quiesce new). Relationship that prevents a system from scheduling jobs that allocate new data sets or modify existing ones in a VIO storage group, a pool storage group, or individual volumes within a pool storage group.

R

read-only variable. An ACS language variable that contains data set or system-derived information. It can be referenced but not altered in an ACS routine.

read-write variable. An ACS language variable that is assigned a value within an ACS routine. It can be referenced, and each ACS routine assigns a value to its own, unique read-write variable.

recall. The process of moving a migrated data set from a level 1 or level 2 volume to a primary volume or to a volume not managed by DFHSM.

Resource Access Control Facility (RACF). An IBM program product that provides access control by identifying and verifying users to the system. RACF authorizes access to resources, logs unauthorized access attempts, and logs accesses to protected data sets.

S

SCDS. See *source control data set.*

secondary allocation. A predefined contiguous space on a DASD volume reserved for additions to a particular data set, and allocated only after the primary allocation space is full.

sequential data set. A data set whose records are organized on the basis of their successive physical positions, such as on magnetic tape.

SMS class. A list of attributes that SMS applies to data sets having similar allocation (data class), performance (storage class), or backup and retention (management class) needs.

SMS complex. A set of up to eight systems within an installation that are defined to MVS in the base configuration as SMS systems.

SMS configuration. A configuration base, SMS class and storage group definitions, and ACS routines that SMS uses to manage storage.

SMS-managed data set. A data set that has been assigned a storage class.

SMS volume. See *CONVERTED*

source control data set (SCDS). A VSAM linear data set containing the SMS class and storage group defi-

nitions, ACS routines, and system information that together form an SMS configuration.

storage administrator. A person in the data processing installation who is responsible for defining, implementing, and maintaining storage management policies.

storage class. A list of storage performance and availability service requests.

storage group. VIO, a list of real DASD volumes, or a list of serial numbers of volumes that no longer reside on a system but that end users continue to reference in their JCL.

Storage Management Subsystem (SMS). An operating environment that helps automate and centralize the management of storage. To manage storage, SMS provides the storage administrator with control over data class, storage class, management class, storage group, and ACS routine definitions.

storage pool. A predefined set of DASD volumes used to store groups of logically related data according to user requirements for service or according to storage management tools and techniques.

synchronization time interval. The number of seconds that SMS allows before it checks the COMMDS for volume status, space statistics, and configurational changes from other systems in the complex.

system-managed storage. An approach to storage management in which the system determines data placement and an automatic data manager handles data backup, movement, space, and security.

system residence (SYSRES) volume. The volume on which the nucleus of the operating system and the master catalog are stored.

T

time sharing option (TSO). An optional configuration of the operating system that provides conversational time sharing from remote stations.

translate. To check ACS routine source code for syntactic and semantic errors. If no errors exist, the

translation process generates an object table from the source code and places the object table into a specified SCDS.

V

valid. A status indicating that an SCDS or ACDS is complete. The ACS routines only assign SMS class names and lists of storage group names that are defined to the SCDS. To be complete, an SCDS must contain at least one storage class definition, at least one pool storage group definition containing at least one volume, a storage group selection routine, and base configuration information. A valid SCDS can be activated.

validate. To check the completeness and consistency of an individual ACS routine or an entire SCDS.

VIO storage group. A type of storage group that allocates data sets to paging storage, which simulates the activity of a DASD volume. VIO storage groups do not contain any actual DASD volumes.

virtual storage access method (VSAM). An access method for direct or sequential processing of fixed and variable length records on DASD.

volume. A certain portion of data, together with its data carrier, that can be mounted on the system as a unit; for example, a tape reel or a disk pack. For DASD, a volume refers to the amount of space accessible by a single actuator.

volume pooling. Grouping volumes and managing them as if they were a single, large volume rather than managing them on an individual basis.

volume serial number. An identification number in a volume label that is assigned when a volume is prepared for use on the system.

volume table of contents (VTOC). A table on a DASD device that describes each data set on the volume.

VSAM volume data set (VVDS). A data set that describes the data set characteristics of VSAM data sets residing on a given volume; part of an integrated catalog facility catalog.

Index

attributes *(continued)*
 pool storage group 31
 retention periods 40
 storage class
  availability 49
  defining 48
 storage group 26
audience xi
AUDIT parameter
 bypass RACF logging 113
automatic class selection routine
 See ACS routine
automatic processing
 pool storage group 31
availability attributes 49

B
background information xi
backup
 AUTO BACKUP 31
 defining management class attributes 43
 management class
  frequency 44
  number of versions 44
  version retention period 44
base configuration
 adding a system 21
 control data set application selection 18
 defining SCDS information 17
 definition 1, 18
 deleting a system 21
 planning SDCS information 18
 renaming a system 21
 SMS complex 1
 starter set 144
  assumptions 143
 systems in the complex 21
bias (storage class) 50
bibliography xii
boolean expressions
 ACS routine 126
browsing
 ACS routine, translation 66, 68

C
CACHE fast write
caching 163
changing
 data class 99
 management class 97
 storage class 98
 storage group 96
command and line operator
 reference summary, ISMF 139
COMMDS (communications data set)
 allocation 10
 control data set sizing 8

COMMDS (communications data set) *(continued)*
 current utilization statistics 8
 definition 8
 JCL allocation 10
 recovery 105
 size calculation 9
 SMS complex 8
communications data set
 See COMMDS
comparison operators
 ACS routine 125
connection
 storage group to volume 35
 system to storage group 28
constructs
 traits and characteristics 1
 types 1
continuous availability (storage class) 50
control data set
 allocating 7
 application selection panel 18
 background 4
 description 7
 fixed data fields 8
 RACF 115
 recovery 105
 sizing formula 8
 using ISMF primary option menu 3
CONVERTED volume status 24
converting volumes to SMS 104
copying SMS classes and storage groups 100
corequisite products 3

D
DASD fast write
data class
 ACS routine
  starter set 158
 altering 99
 assigning 60
 attributes
  applying to RECORG 60
  description 54
  processing in JCL 60
 creating additional 61
 defining volume 57—60
 definition
  application selection 54
  record and space attributes 55—57
 deleting 103
 description 53
 displaying list 94
 order when processing attributes 60
 planning 53
 record and space attributes 55
 specifying sort criteria 94
 starter set 149
 volume and VSAM attributes 57

MVS/ESA
Storage Administration Reference

SC26-4514-2

This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you. Your comments will be sent to the author's department for whatever review and action, if any, are deemed appropriate.

Note: Do not use this form to request IBM publications. If you do, your order will be delayed because publications are not stocked at the address printed on the reverse side. Instead, you should direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

If you have applied any technical newsletters (TNLs) to this book, please list them here: _____

Comments (please include specific chapter and page references) :

If you want a reply, please complete the following information:

Name _____ Date _____

Company _____ Phone No. (_____) _____

Address _____

Thank you for your cooperation. No postage is 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 them directly to the address in the Edition Notice on the back of the title page.)

SC26-4514-2

Reader's Comment Form

IBM
®

MVS/ESA
Storage Administration Reference

SC26-4514-2

This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the under-standing that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you. Your comments will be sent to the author's department for whatever review and action, if any, are deemed appro-priate.

Note: Do not use this form to request IBM publications. If you do, your order will be delayed because publications are not stocked at the address printed on the reverse side. Instead, you should direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

If you have applied any technical newsletters (TNLs) to this book, please list them here: _____

Comments (please include specific chapter and page references) :

If you want a reply, please complete the following information:

Name _____   Date _____

Company _____   Phone No. (_____) _____

Address _____

_____

Thank you for your cooperation. No postage is 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 them directly to the address in the Edition Notice on the back of the title page.)

SC26-4514-2

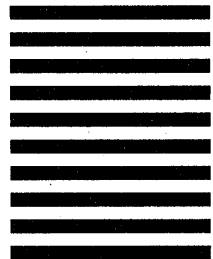**Reader's Comment Form**

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 40 ARMONK, NY

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department J57
P.O. Box 49023
San Jose, CA 95161-9945

IBM ®