

STORAGE MANAGEMENT AT ITS

CAPACITY AND PERFORMANCE MANAGEMENT GROUP

INFORMATION TECHNOLOGY SERVICES

STATE OF NORTH CAROLINA

CLIENT USERS GUIDE




TABLE OF CONTENTS

DATA SETS	3
ALLOCATING DATASETS	3
<i>METHOD 1—ISPF UTILITIES MENU</i>	4
NAMING DATASETS	6
DATASET NAMING STANDARDS	6
DISK MANAGEMENT	9
SYSTEM MANAGED STORAGE (SMS)	9
DATA CLASS	10
DATA CLASS ATTRIBUTES	12
DATA CLASS ATTRIBUTE DEFINITIONS	15
STORAGE GROUP	22
STORAGE GROUP DESCRIPTIONS	22
STORAGE GROUP SERVICES	27
STORAGE GROUP SERVICE DEFINITIONS	29
STORAGE CLASS	29
MANAGEMENT CLASS	31
MANAGEMENT CLASS ATTRIBUTES	32
MANAGEMENT CLASS ATTRIBUTE DEFINITIONS	34
EXAMPLES OF DATASET NAMES	37
DB2 DATASETS	37
<i>TEST</i>	37
<i>PRODUCTION</i>	37
TSO DATASETS	39
GENERAL DATASETS	39
<i>PRODUCTION</i>	39
<i>TEST</i>	40
HIERARCHICAL STORAGE MANAGER (HSM)	41
HSM COMMANDS	41
TAPE MANAGEMENT	43

LIST OF TABLES

Table 1	Data Class Attributes, Part 1	12
Table 2	Data Class Attributes, Part 2	13
Table 3	Data Class Attributes, Part 3	14
Table 4	Storage Group Services	27
Table 5	Storage Class and corresponding Storage Group	30
Table 6	Management Class Attributes, Part 1	32
Table 7	Management Class Attributes, Part 2	33

STORAGE MANAGEMENT AT ITS

This document provides information about storage management policies at ITS. It is intended to be a basic guide to storage management for ITS clients who manage their own data using the IBM Storage Management Subsystem (SMS) software product. Following the recommendations in this guide can enable SMS to provide time and cost saving advantages including:

- Automatic determination of many dataset allocation parameters
- Automatic selection of appropriate storage devices
- Automatic management of dataset retention and backup
- Reduction in DASD space used by datasets

Most of the dataset allocation parameters selected by SMS are based on the name assigned to the dataset when it is created. This guide therefore details the ITS Standard Dataset Naming Conventions.

Examples of methods for creating datasets are included for both batch jobs and interactively from ISPF.

A complete SMS guide is beyond the scope of this document. This guide addresses how datasets are named, allocated, and managed at ITS using System Managed Storage (SMS), and Hierarchical Storage Management (HSM). Complete IBM documentation of HSM and SMS is available at ITS from TSO using Book Manager from ISPF.

Questions regarding this document should be directed to the Storage Management team in ITS Capacity and Performance Group.

TERMS USED IN THIS DOCUMENT

ACS	Automatic Class Selection
ACTS	Automated Collection and Tracking System
ANSI	American National Standards Institute
BCDS	Backup Control Dataset
CICS	Customer Information Control System
CS	Computing Services, a division of ITS
DASD	Direct Access Storage Device
DB2	Database 2
DBA	Database Administrator
DFSMSHsm	Data Facility Storage Management Subsystem hierarchical storage manager
DHHS	Department of Health and Human Services
DOR	Department of Revenue
DOT	Department of Transportation
DS	Dataset
EIS	Eligibility Information System
ES	VSAM Entry Sequenced Dataset
EXPDT	Expiration Date
GDG	Generation Data Group
GDS	Generation Dataset
HLQ	High Level Qualifier
HSM	Hierarchical Storage Manager
IMS	Information Management System
ISPF	Interactive System Productivity Facility
ISPF/PDF	ISPF /Program Development Facility
ITS	Information Technology Services
JCL	Job Control Language
KS	VSAM Keyed Sequential Dataset
KSDS	Key Sequential Dataset
LLQ	Low Level Qualifier

LPAR	Logical Partition
LS	VSAM Linear Space Dataset
MCDS	Migration Control Dataset
ODS	Output Dataset
PDS	Partitioned Dataset
PDSE	Partitioned Dataset Extended
PO	Partitioned Organization
PS	Physical Sequential
RETPD	Retention Period
RLS	Record Level Sharing
RMDS	Report Management and Distribution System
RR	VSAM Relative Record Dataset
SMS	System Managed Storage
STARS	State Title and Registration System
TSO	Time Sharing Option
VSAM	Virtual Storage Access Method

DATA SETS

ALLOCATING DATASETS

The process of creating a new dataset is called "allocating" a dataset. This process accomplishes two tasks:

1. Reserves disk space so there is a place for the information that the dataset will contain; and
2. Enters the name and volume serial of the dataset in the system catalog so that the operating system can find the dataset the next time you request it.

Most datasets are created by batch processing. Automatic Class Selection (ACS) routines in SMS have been configured by CS to assign appropriate class information to your datasets based on the allocation parameters and the dataset name specified in your JCL. It is best to review the current data classes that have been defined based on common allocation parameters. By specifying the data class, you no longer need to code the Volume serial number, device type, DCB parameters, or even the amount of space in your JCL. If circumstances arise that require different allocation parameters, they can still be specified in JCL to override the ACS assignments.

When blocksize is unspecified or set to zero in your JCL, SMS will determine the most efficient blocksize based on the record length, dataset organization, and device type. This can greatly improve utilization of DASD, decreasing the total cost of data storage and frees the user from complex calculations based on device geometry to determine an efficient blocksize for datasets.

ACS routines override JCL specified Management Class and Storage Class

```
//SMSDS2 DD DSNAME=MYDS2.PGM,STORCLAS=SCLAS02,DISP=(NEW,KEEP)
//MYDATA DD DSN=datasetname, SPACE=(TRK,(1,1)),
//   MGMTCLAS=STANDARD,,STORCLAS=BASE,
//   RECFM=FB,LRECL=307,BLKSIZE=0,DISP=(,CATLG,DELETE)
```

Method 1—ISPF UTILITIES MENU

- Step 1. To allocate a dataset, at the *ISPF Primary Option Menu* select option 3, **Utilities**. This will bring up the *Utility Selection Menu*.
- Step 2. At the *Utility Selection Menu*, select option 2, **Dataset**. This will bring up the *Data Set Utility* screen.
- Step 3. At the *Data Set Utility* screen, select A, **Allocate new data set**. Press the **TAB** key to move to **Data Set Name**, enter the dataset name (see *Naming Datasets* below) and press the **Enter** key. This will bring up the *Allocate New Data Set* screen.
- Step 4. The *Allocate New Data Set* screen contains a list of parameters for the new dataset. Most of the parameters will already contain information copied from the last dataset you created. Change the parameters by pressing the **Tab** key to move to that parameter and typing over the existing entry. Make sure that the **Volume Serial** parameter is left blank to create an SMS dataset. Press the **Enter** key to accept the parameters for the dataset. This will return

you to the *Data Set Utility* screen. The upper right-hand corner of the screen will contain the message:

DATA SET ALLOCATED

In order to allocate a dataset and make full use of storage management options at ITS, it is essential to understand dataset naming conventions and SMS. The following sections will cover dataset naming standards and the ACS Routines established for Data Class, Storage Class, Management Class, and Storage Groups. Carefully review the attributes for each Data Class, Storage Class, Management Class, and Storage Group at ITS.

NAMING DATASETS

The dataset name is assigned to a particular set of information and distinguishes that set of information from others on the disk or tape. Dataset names must conform to specific standards established at ITS.

DATASET NAMING STANDARDS

Standard 1 ~ All Datasets

The operating system requires that all dataset names are strings of up to 44 characters that conform to the following restrictions:

- Characters are alphabetic (A-Z), numeric (0-9), national (@, #, and \$) and periods (.)
- The whole name must be divided into segments (qualifiers) of up to eight characters each. For example, SEPT.REPORT.DATA is the name of a dataset whose qualifiers are SEPT, REPORT, and DATA.
- The first character of the dataset name and the first character after a period must be alphabetic or national
- The last character of a dataset name may **not** be a period, nor may there be two consecutive periods.

Standard 2 ~ Non-DB2 and Non-TSO datasets

Non-DB2 and non-TSO dataset names require two qualifiers.

- The high-level qualifier must be the client's three-letter department code.
- The second-level qualifier must be the client's three letter application code.
- The third-level qualifier should consist of the following variables to create a six-character qualifier:

PROD	Production dataset or
or	
TEST	Test dataset
O	Datasets allocated by CICS or IMS
or	
B	All other datasets
I or	DLI
F or	FOCUS
L or	LIBRARY
P or	PDS
or V	VSAM
or A	SAS,
or S	SEQ
or Y	SYSTEM

For example, a non-TSO or non-DB2 dataset name would have the following format:

AAA.BBB.PRODO I.*.*

In this example AAA is the client's three letter department code, and BBB is the client's three letter application code, PROD indicates a production dataset, O indicates a dataset allocated by CICS or IMS, I indicates a DLI dataset, and the asterisks indicate a client selected LLQ.

Following the first three qualifiers, subsequent levels are specified by the client. When you are naming datasets it is crucial to think about storage management considerations. The qualifiers used in the dataset name regulate how datasets are allocated, stored and managed at ITS. ACS Routines are the rules that determine the data class, storage class, storage group, and management class for the dataset.

Standard 3 ~ TSO datasets

TSO datasets have three qualifiers.

- The high-level qualifier consists of the department code and TSO
- The second-level qualifier consists of TEAM and the SCC assigned team number.
- The third-level qualifier consists of a client selected LLQ.

For example, a TSO dataset name would have the following format:

aaaTSO.TEAMbb.ccccccc

In this example, *aaa* is the department code, *bb* is the SCC assigned team number, and *ccccccc* is the LLQ.

Standard 4 ~ DB2 datasets

If VSAM datasets are to be explicitly defined for table spaces, DB2 requires a special naming convention. Each dataset is composed of the following qualifiers:

aaaDBxxx	VCAT Alias; limited to 8 characters. First three characters must be a valid ITS department code, followed by DBxxx, the DB2 subsystem ID.
DSNDBC or DSNDBD	The first five characters must be DSNDB, followed by C for VSAM clusters or D for VSAM data components
bbbDxxxx	Database Name. First three characters must be a valid application billing code, followed by D (for database), and any four characters.
cccSxxxx cccXxxxx	Table Space or Index Space Name. First three characters must be a valid application billing code followed by S for Table Space or X for Index Space, and any four characters.
I0001	Constant Indication, use for all dataset names
A001	Standard – must be keyed as shown if table space is not partitioned

For example, a DB2 dataset name would have the following format:

aaaDBxxx.DSNDBD.bbbDxxxx.cccSxxxx.I0001.A001 Table Space

aaaDBxxx.DSNDBD.bbbDxxxx.cccXxxxx.I0001.A001 Index Space

In these examples *aaaDBxxx* is an alias consisting of the valid ITS department code, and the DB2 subsystem ID, DSNDBD indicates a VSAM data component, *bbbDxxxx* is the database name, *cccSxxxx* for the Table Space name or *cccXxxxx* for the Index Space name, *I0001* is the Constant Indication, and *A001* indicates that the table space is not partitioned.

Standard 5 ~ SYSW LPAR datasets

Dataset names on the SYSW LPAR should use the naming convention AAAT or AAAY for the first high-level qualifier, where AAA is the three-digit department code, T represents a test and development file, and Y represents a Year 2000 test and development file.

Standard 6 ~ HFS / ZFS datasets

In z/OS, UNIX System Services (USS) files and directories are stored in Hierarchical File System (HFS) & zSeries File System (ZFS) data sets. Like MVS data sets, HFS/ZFS data sets can be allocated, backed up, restored, and deleted. HFS/ZFS datasets must be mounted at specific points in the UNIX directory tree structure before they can be accessed through USS. Mountpoints can be requested by contacting our Customer Support Center at 754-6000, or toll free, 1-800-722-3946. A Remedy ticket should be requested and assigned to General Systems.

The following naming standards have been established for HFS/ZFS files:

xxx.yyy.**.SYSz.HFS

xxx.yyy.**.SYSz.ZFS

xxx The high-level qualifier must be the client's three letter department code, as shown on the invoice. An additional T indicates the dataset is for test.

yyy The second-level qualifier must be the client's three letter application code.

** The next level qualifier is freeform. The total length cannot exceed 30 chars.

SYSz.HFS/ZFS The dataset name must end in SYSz.HFS/ZFS. Where z is the LPAR, which the dataset is mounted.

The following are examples of valid dataset names:

DHR.BDA.SYSA.HFS / DHR.BDA.SYSA.ZHS

OSCT.AFR.NCAS.HOD.SYSW.HFS

DISK MANAGEMENT

DASD management at ITS ensures integrity of your data on disks, maintains maximum disk space availability, and promotes optimum use of disk space. The DASD management practices that most directly affect you are those related to space release, dataset migration (archiving), and the daily backup of new or modified disk datasets.

SYSTEM MANAGED STORAGE (SMS)

System Managed Storage allows the operating system to take over many storage management tasks that previously were performed manually. The purpose of SMS is to eliminate, simplify and automate tasks normally done by users or a storage administrator. SMS uses software programs to manage data security, placement, migration, backup, recall, recovery, and deletion of datasets. This ensures that current data is available when needed, and obsolete data is removed from storage.

The goals of system-managed storage are to:

- Reduce computing costs by improving the use of the storage media; for example, by reducing out-of-space ABENDs and providing a way to set a free-space requirement.
- Reduce labor costs involved in storage management by centralizing control, automating tasks, and providing interactive controls for storage administrators. This allows the clients to focus on their job responsibilities.
- Reduce the client's need to be concerned with the physical details of performance, space, and device management. Clients can focus on using data instead of managing data.

SMS is tailored to meet individual needs. SMS allows the storage administrator to affect new dataset allocations, volume extend, and perform data movement by assigning classes and groups to a new dataset, which affect dataset attributes, placement, migration, and backup.

Classes and groups are assigned via installation-written ACS Routines. The assigned classes and groups are:

- used by SMS to determine dataset characteristics during initial allocation
- used by SMS to select the best available volume(s) for dataset allocation
- used by SMS for dynamic cache management
- used by HSM for backup and migration decisions
- used by HSM for backup and recovery

All data, storage, and management classes can be viewed via option O of the *ISPF Primary Option Menu*, option O of "*OTHER CMDS*", and option S of "*SYSCMDS*".

The type of data you are storing determines the data class that you select. The next section describes the data classes already set up at ITS, and rules in the ACS Routines that assign datasets to these data classes. Before allocating datasets, review the attributes for the data classes and the ACS Routines that assign datasets to each Data Class.

DATA CLASS

Data Classes are used to allocate attributes for a dataset in one operation. Data classes apply to SMS-managed and non-SMS datasets. They allow you to define allocation defaults and simplify allocations by using data class in place of many keywords. In order to allocate the attributes for the dataset, the data class name must be explicitly specified in the client's JCL or implicitly via the ACS Routines. SMS uses only those data class attributes that have meaning for the dataset being allocated. All Data Class parameters (with the exception of DATASET NAME TYPE) may be overridden with JCL or other allocation methods. CS will implicitly assign the following data classes via the ACS Routines:

DATA CLASS NAME	Rule that assigns the Data Class to the dataset
DATAF	<p>When the third qualifier of the dataset name is any of the following:</p> <p>PRODOF PRODOV PRODOS PRODBF PRODBP PRODBA PRODBS TESTOF TESTOP TESTOS TESTBF TESTBP TESTBS</p> <p>OR the last qualifier is one of the following:</p> <p>DATA F DATA.</p>
DATAV	<p>When the last qualifier of the dataset name is one of the following:</p> <p>TEXT VDAT A.</p>
LOADLIB	<p>When the last qualifier of the dataset name is one of the following:</p> <p>RESLIB PGMADB LINKLIB LOAD*</p> <p>OR the third qualifier is one of the following:</p> <p>PRODBL PRODBY TESTOL TESTOY TESTBL TESTBY.</p>
SRCFLIB	<p>When the last qualifier of the dataset name is one of the following:</p> <p>COB* ASM* CNTL* PROCLIB SCR, SOURCE FOR* JCL*.</p>
SRCVLIB	<p>When the last qualifier is one of the following:</p> <p>PLI PL1 SCRIPT CLIST.</p>
LISTING	<p>When the last qualifier is one of the following:</p> <p>SYSOUT OUTLIST LIST LISTING.</p>
KEYED	<p>When the last qualifier is one of the following:</p> <p>PRODOI PRODBI PRODBV TESTOI TESTOV TESTBI TESTBV.</p>
DEFAULT	<p>This is the default data class. This data class is assigned to all datasets except where the specific qualifiers are targeted for other data class names or the data class name has been specified in the JCL.</p>

DATA CLASS ATTRIBUTES

The following three tables list data classes and their attributes that are available for client specification. Definitions and possible values follow in *Data Class Attribute Definitions*.

Table 1 Data Class Attributes, Part 1

ATTRIBUTES	DATA CLASS NAME					
	DATAF	DATAV	DIRECT	ENTRY	KEYED	LINEAR
RECORG	--	--	RR	ES	KS	LS
RECFM	FB	VB	--	--	--	--
LRECL	80	225	--	--	--	--
KEYLEN	--	--	--	--	--	--
KEYOFF	--	--	--	--	0	--
AVGREC	U	U	U	U	U	U
AVG VALUE	80	255	4,096	4,096	4,096	4,096
SPACE PRIMAR	5,000	5,000	100	100	100	100
SPACE SECON	5,000	5,000	100	100	100	100
SPACE DIRECTORY	---	---	---	---	---	---
RETPD or EXPDT	---	---	---	---	---	---
ADDITIONAL VOLUME AMT	---	---	---	---	---	---

CISIZE DATA	---	---	---	---	---	---
%FREE SPACE CA	---	---	---	---	10	---
%FREE SPACE CI	---	---	---	---	10	---
SHARE XREGION	---	---	---	---	2	---
SHARE XSYSTEM	---	---	---	---	3	---
DATASET NAME TYPE	---	---	---	---	---	---
EXTD ADDRESS.	---	---	---	---	---	---
REUSE	---	---	---	---	---	---
INITIAL LOAD				---	---	---
SPANNED/NONSPANNED	---	---	---	---	---	---
BWO	---	---	---	---	---	---
LOG	---	---	---	---	---	---
LOGSTREAM ID	---	---	---	---	---	---
REC ACC BIAS	USER	---	---	---	---	---

Table 2 Data Class Attributes, Part 2

ATTRIBUTES	DATA CLASS NAME					
	LISTING	LOADLIB	SRCFLIB	SRCVLIB	DEFAULT	TEST
RECORG	--	--	--	--	--	--
RECFM	VBA	U	FB	VB	--	--
LRECL	137	---	80	255	--	--
KEYLEN	---	---	---	---	---	---
KEYOFF	---	---	---	---	---	---
AVGREC	U	U	U	U	U	---
AVG VALUE	137	23,476	80	255	---	---
SPACE PRIM C	20,000	50	5,000	5,000	5,000	---
SPACE SECON	20,000	50	5,000	5,000	5,000	---
SPACE DIRECTORY	----	62	62	62	----	----
RETPD or EXPDT	----	----	----	----	----	----
ADDITIONAL VOLUME AMT	---	---	---	---	---	---
CSIZE DATA	----	----	----	----	----	----
%FREE SPACE CA	----	----	----	----	----	----
%FREE SPACE CI	----	----	----	----	----	----
SHARE XREGION	----	----	----	----	----	----
DATASET NAME TYPE	---	---	---	---	---	---
SHARE XSYSTEM	----	----	----	----	----	----
EXTD ADDRESS.	YES	---	---	---	---	---
REUSE	NO	NO	NO	---	---	---
INITIAL LOAD				---	---	---
SPANNED/NONSPANN ED	---	---	---	---	---	---
BWO	---	---	---	---	---	---
LOG	---	---	---	---	---	---
LOGSTREAM ID	---	---	---	---	---	---
REC ACC BIAS	USER	---	---	---	---	---

Table 3 Data Class Attributes, Part 3

ATTRIBUTE	DATA CLASS NAME		
	EXFKSDS	KEY4096	SEQ4096
RECORG	KS	KS	ES
RECFM	---	---	---
LRECL	---	---	---
KEYLEN	---	---	---
KEYOFF	---	---	---
AVGREC	U	U	U
AVG VALUE	4096	4096	4096
SPACE PRIM	100	100	100
SPACE SECON	100	100	100
SPACE DIRECT	----	---	---
RETPD or EXPDT	----	----	----
ADDITIONAL VOLUME AMT	---	---	---
CISIZE DATA	----	4096	4096
%FREE SPACE CA	10	10	----
%FREE SPACE CI	10	10	----
SHARE XREGION	2	2	2
SHARE XSYSTEM	3	3	3
DATASET NAME TYPE	---	---	---
EXTD ADDRESS.	YES	---	---
REUSE	NO	NO	NO
INITIAL LOAD	RECOVERY	RECOVERY	RECOVERY
SPANNED/ NONSPANNED	---	---	---
BWO	---	---	---
LOG	---	---	---
LOGSTREAM ID	---	---	---
REC ACC BIAS	USER	---	---

DATA CLASS ATTRIBUTE DEFINITIONS

RECORDG

The RECORDG field shows how VSAM datasets allocated by a Data Class are organized.

Possible values:

KS	VSAM Keyed Sequential Dataset
ES	VSAM Entry Sequenced Dataset
RR	VSAM Relative Record Dataset
LS	VSAM Linear Space Dataset

A blank field (the default) means that the Data Class is used for non-VSAM datasets having Partitioned Organization (PO) or Physical Sequential (PS) organization.

RECFM

The RECFM field shows the record format and type of carriage control assigned to non-VSAM datasets:

Possible values (they may appear singly or in combination):

B	Blocked format
F	Fixed format
S	Standard (if format fixed) or Spanned (if format is variable)
U	Undefined format
V	Variable format
A	ANSI carriage control
M	Machine carriage control

A blank field (the default) means that the Data Class is used for non-VSAM datasets.

LRECL

The LRECL field shows, in bytes, the logical record length. The value shown is the length of fixed-length records or the maximum length of variable-length records.

Possible values:

1 to 32760 or blank	for non-VSAM datasets
1 to 32761 or blank	for VSAM datasets

	If you are using a key-sequenced VSAM dataset, the KEYLEN attribute cannot be larger than the logical record length.
KEYLEN	<p>The KEYLEN field shows, in bytes, the size of each record key in a non-VSAM dataset, or the size of each key field in a key-sequenced VSAM dataset.</p> <p>Possible values:</p> <p>0 to 255 for non VSAM datasets.</p> <p>1 to 255 for key sequenced VSAM datasets</p>
KEYOFF	<p>The KEYOFF field applies only to key-sequenced VSAM datasets. The field shows, in bytes, the distance from the start of the record to the start of the key field.</p> <p>Possible values:</p> <p>0 to 32760</p> <p>KEYOFF must be in the range 0 to the difference between LRECL and KEYLEN.</p>
AVGREC	<p>The AVGREC field shows whether this Data Class allocates space in bytes, kilobytes, or megabytes.</p> <p>Possible values:</p> <p>K Space is allocated in kilobytes</p> <p>M Space is allocated in megabytes</p> <p>U Space is allocated in bytes.</p>
AVG VALUE	<p>The AVG VALUE field shows the multiplication factor used in determining allocated space. Primary space equals AVG VALUE times SPACE PRIMARY; secondary space equals AVG VALUE times SPACE SECONDARY. In both cases, space is allocated in bytes if AVGREC = U, in kilobytes if AVGREC = K, or in megabytes if AVGREC = M. Thus, if AVG VALUE = 80, SPACE PRIMARY = 1000, and AVGREC = U, it means that 80,000 bytes of primary space have been allocated.</p> <p>Possible values:</p> <p>0 to 65535</p>
SPACE PRIMARY	<p>The SPACE PRIMARY value; when multiplied by AVG VALUE, determines the amount of space that this Data Class initially allocates for a dataset. Space will be allocated in bytes if AVGREC = U, in kilobytes if AVGREC = K, or in megabytes if AVGREC = M. Thus if SPACE PRIMARY = 200, AVG VALUE = 1, and AVGREC = K, the initial space allocated will be 200K (200 X 1).</p> <p>Possible values:</p> <p>0 to 999999</p>
SPACE SECONDARY	<p>The SPACE SECONDARY value, when multiplied by AVG VALUE, determines the additional space that can be allocated for a dataset. Space will be allocated in bytes if AVGREC = U, in kilobytes if AVGREC = K, or in megabytes if AVGREC = M. Thus if SPACE SECONDARY =</p>

50, AVG VALUE = 6, and AVGREC = M, the additional space allocated will be 300 (50 X 6) megabytes.

Possible values:

0 to 999999

SPACE DIRECTORY

The SPACE DIRECTORY field shows the number of blocks allocated for the directory of a partitioned dataset.

Possible values: 0 to 999999

RETPD OR EXPDT

The RETPD OR EXPDT field shows either the retention period or the expiration date assigned to datasets by this Data Class. Datasets will be deleted or archived either one day after the retention period or on the expiration date. You can override the Data Class value with JCL or other allocation methods, but they don't appear in the data column.

Possible values:

yyyy/mm/dd Any year from 1900 to 2155; month: 01 to 12; day: 01 to 31

yyyy/00/00 The expiration date has been set to this special value, which is meaningful to other programs. The year yyyy is in the range 1900-2155

0 to 9999 Datasets expire in the number of days shown

The following special characters can appear in this column:

----- No retention period or expiration date has been specified for this data class

???????????? The value cannot be displayed because of a data conversion error

ADDITIONAL VOLUME

The ADDITIONAL VOLUME AMT field shows the type of allocation amount when a VSAM dataset in extended format begins allocation on subsequent new volumes.

Possible values:

PRIMARY Primary allocation amount has been requested.

SECONDARY Secondary allocation amount has been requested. If the value has not been specified. The system will use the default value of PRIMARY.

CISIZE DATA

The CISIZE DATA field is for VSAM datasets with a RECORG of ES, KS, LS, or RR. This field shows the number of bytes allocated for each control interval in data portion, not the index portion, of a dataset.

To allow for overhead processing, the CISIZE value will be at least seven bytes greater than maximum record size.

Possible values:

1 to 32768

%FREE SPACE CA

The %FREE SPACE CA field shows what percentage of each control area in a key-sequenced VSAM dataset should be set aside as free space. VSAM uses the space to lengthen or insert records, as needed.

Possible values:

0 to 100

%FREE SPACE CI

The %FREE SPACE CI field shows what percentage of each control interval in a key-sequenced VSAM dataset should be set aside as free space. VSAM uses the space to lengthen or insert records, as needed.

Possible values:

0 to 100

SHARE XREGION

The SHARE XREGION field shows how a VSAM dataset can be shared among regions of one system, or across regions of multiple systems.

Possible values:

1 All users can read the dataset OR one user can update it.

2 All users can read the dataset AND one user can update it.

	3	All users can read and update the dataset. VSAM does not protect the dataset
	4	All users can read and update the dataset. VSAM helps prevent lost, damaged, or altered data.
SHARE XSYSTEM		The SHARE XSYSTEM field shows how a VSAM dataset can be shared among systems.
	3	All users can read and update the dataset. VSAM does not protect the dataset.
	4	All users can read and update the dataset. VSAM monitors access of datasets in order to prevent lost, damaged, or altered data.
DATASET NAME TYPE		This column shows the format that is used to allocate datasets using this data class. Possible values:
	EXTENDED PREFERRED	The system allocates datasets in extended format if the necessary system resources are available. Otherwise, it allocates them in non-extended format.
	EXTENDED REQUIRED	The system allocates data sets in extended format if the necessary system resources are available. Otherwise, the allocations fail.
	LIBRARY	The system allocates datasets as PDSEs.
	PDS	The system allocates datasets as PDSs.
EXTENDED ADDRESSABILITY		The extended addressability field specifies whether or not a VSAM KSDS in the extended format is allowed to grow beyond the four gigabyte (4GB) size. Possible values:
	YES	Provides extended addressability.
	NO	Does not provide extended addressability.
REUSE		The REUSE field indicates whether or not users can open the cluster again and again as a new cluster. Possible values:
	YES	Indicates that the cluster is reusable
	NO	Indicates that the cluster is not reusable

INITIAL LOAD

The INITIAL LOAD field indicates whether or not the storage allocated to the data component was preformatted before records were inserted during initial load.

Possible values:

SPEED	Indicates that the data component's space was not preformatted
RECOVERY	Indicates that the data component's space was preformatted

**SPANNED/
NONSPANNED**

The SPANNED/NONSPANNED field shows whether the data record is allowed to cross control interval boundaries.

Possible values:

SPANNED	Indicates that the record is contained in more than one control interval.
NONSPANNED	Indicates that the record is contained in one control interval

The following special characters may appear in this data column:

-----	If the value has not been specified.
???????????	Value cannot be displayed because of data conversion error.

BWO

The BWO field shows that backup-while-open (BWO) is allowed for the sphere.

Possible values:

TYPECICS	BWO processing for CICS VSAM file control datasets is allowed
TYPE IMS	BWO processing for IMS VSAM datasets is to be used
NO	BWO does not apply to this cluster

The following special characters may appear in this data column:

-----	If the value has not been specified
???????????	Value cannot be displayed because of a data conversion error

LOG

The LOG field shows whether the sphere to be accessed with VSAM RLS is recoverable or non-recoverable.

Possible values:

NONE	Not a recoverable dataset.
UNDO	Changes to the sphere accessed in RLS mode can be backed out.
ALL	Both the external backout and the forward recovery capability are available for the sphere accessed in RLS mode.

LOGSTREAM ID

The LOGSTREAM ID shows the name of the recovery logstream.

Possible values:

Dataset name Prefix name for the recovery logstream dataset.

REC ACC BIAS

The RECORD ACCESS BIAS field allows the system to acquire and choose the buffering algorithms.

Possible values :

SYSTEM	A value of SYSTEM allows the system to choose the number of buffers and the buffering algorithms for the VSAM dataset.
USER	A value of USER prevents SYSTEM MANAGED BUFFERING and number of buffers and buffering algorithms will be based on user specified (or defaulted) values and current algorithms.

STORAGE GROUP

The Storage Class selected will determine the Storage Group for the dataset. Storage groups separate data by client or by type. If you look at Table 5 you will see the corresponding Storage Groups and Storage Classes.

Storage groups are a collection or pool of storage volumes that are defined to SMS. SMS allows the use of various management classes and various storage classes in a single SMS storage group. Storage groups reduce the need for the user to understand the physical requirements of the storage devices. SMS allows users to specify a data, storage, or management class. The Storage Group that is assigned is dependent upon the Storage Class that is assigned.

STORAGE GROUP DESCRIPTIONS

Please read through the descriptions of the Storage Groups to determine the Storage Group that holds the type of data you are allocating. After determining the best Storage Group, go to *Storage Class* and review Table 5.

STORAGE GROUP NAME	STORAGE GROUP DESCRIPTION
ACTSDB	This storage group contains the DHHS ACTS production DB2 databases. The client department manages this storage group, with support provided by CS. For the specific database naming conventions contact your DBA.
ACTSDBRO	This storage group contains the DHHS ACTS production read-only DB2 databases. The client department manages this storage group, with support provided by CS. For the specific database naming conventions contact your DBA.
ACTSTEST	This storage group contains the DHHS ACTS test DB2 databases. For specific test database naming conventions contact your DBA. The client department manages this storage group, with support provided by CS.
ACTSVRU	This storage group contains the DHHS ACTS production DB2 databases to support Voice Response Unit child support information requests. For specific database naming conventions contact your DBA. The client department manages this storage group, with support provided by CS.

STORAGE GROUP NAME	STORAGE GROUP DESCRIPTION
CORPERF2	This storage group contains specifically managed COR datasets which must be isolated based on their extremely heavy usage. This improves the performance to those datasets and reduces their heavy usage impact upon other datasets. The data requires special placement on specific volumes. To explore this option, contact the CS Capacity and Performance staff. CS manages this storage group.
DB2PROD	This storage group contains client production DB2 databases, with exceptions for large DB2 applications where the client department has chosen to manage a private storage group. Segregation of the DB2 applications helps ensure data integrity and consistent performance. CS targets this unique data to this storage group. Also, databases that are allocated specifying a storage class of DBBASE will be targeted to this storage group. The non-DB2 databases currently allocated in this storage group are being reallocated to the GENPROD storage group. The test DB2 databases with storage classes DBTEST and DLDBTEST will be reallocated into storage group DB2TEST. CS manages this storage group.
DB2TEST	This storage group contains the client test DB2 databases. For specific database naming conventions contact your DBA. CS manages this storage group.
DLDB90	This storage group contains the DOT Driver License production DB2 databases. This storage group is managed by the client department, with support provided by CS. For the specific database naming conventions contact your DBA.
EISTEST	This storage group contains the DHHS EIS test VSAM and other general datasets. For the specific dataset naming convention contact your application manager. The client department manages this storage group, with support provided by CS.
EISTIMS1	This storage group contains the DHHS EIS test IMS databases. For specific test database naming conventions contact your DBA. The client department manages this storage group, with support provided by CS.
EISTIMS2	This storage group contains the DHHS EIS test IMS databases. For specific test database naming conventions contact your DBA. The client department manages this storage group, with support provided by CS.

STORAGE GROUP NAME	STORAGE GROUP DESCRIPTION

GENPROD

This storage group contains the SMS-managed general production data that has been targeted by the storage administrator. General production data includes interactive data such as TSO and ISPF/PDF libraries, online files or databases, batch files, and VSAM not being targeted to other specific storage groups. DB2 databases should not be targeted into this pool. Datasets that have not been identified as SMS-managed, but specify a storage class of BASE or TSOBASE during allocation will also be allocated in this storage group.

STORAGE GROUP NAME	STORAGE GROUP DESCRIPTION
GENTEST	<p>This storage group contains general test data. These are datasets with the naming convention</p> <p>%%%T.** or %%%TSOT.**</p> <p>or if one of the qualifiers is one of the following:</p> <p>TEST* DEV E*.</p> <p>By contacting CS, datasets with other specific identification as test data maybe requested for inclusion. CS manages this storage group.</p>
LARGE	<p>This storage group contains specifically managed datasets based on their size. Currently, datasets greater than or equal to approximately 2000 cylinders have been targeted by name to this storage group. To explore this option, contact the CS Capacity and Performance staff. CS manages this storage group.</p>
LARGET	<p>This storage group has the same criteria as LARGE except it is for test datasets. CS manages this storage group.</p>
PERFORM	<p>This storage group contains specifically managed datasets which must be isolated based on their extremely heavy usage. This improves the performance to those datasets and reduces their heavy usage impact upon other datasets. The data requires special placement on specific volumes. To explore this option, contact the CS Capacity and Performance staff. CS manages this storage group.</p>
REVENUE	<p>This storage group contains some general datasets and DB2 databases for the DOR. The client department manages this storage group, with support provided by CS. For the specific naming conventions contact your department representative.</p>
REVENUEW	<p>This storage group contains some general test datasets, and DB2 test databases for the Department DOR. The client department manages this storage group, with support provided by CS. For the specific naming conventions contact your department representative.</p>

STORAGE GROUP NAME	STORAGE GROUP DESCRIPTION
-----------------------	---------------------------

VRDB90	This storage group contains the DOT STARS production DB2 databases for the vehicle registration application. For specific database naming conventions contact your DBA. The client department manages this storage group, with support provided by CS.
---------------	--

STORAGE GROUP SERVICES

The following table reflects various services for the storage groups that would be of interest to the clients. Definitions and possible values follow in *Storage Group Service Definitions*.

Table 4 Storage Group Services, Part 1

STORAGE GROUP NAME	SERVICES		
	Automatic Migration	Automatic Backup	Automatic Dump
ACTSDB	NO	NO	NO
ACTSDBRO	NO	NO	NO
ACTSTEST	YES	YES	NO
ACTSVRU	NO	NO	NO
CORPERF2	NO	NO	NO
DB2PROD	NO	NO	NO
DB2TEST	NO	NO	NO
DHRPERF	NO	NO	NO
DLDB90	NO	NO	NO
EISTEST	PRIMARY	NO	NO
EISTIMS1	NO	NO	NO
EISTIMS2	NO	NO	NO
GENPROD	YES	YES	NO
GENTEST	YES	YES	NO

Table 4 Storage Group Services, Part 2

STORAGE GROUP NAME	SERVICES		
	Automatic Migration	Automatic Backup	Automatic Dump
LARGE	Primary	YES	NO
LARGET	Primary	YES	NO
PERFORM	NO	NO	NO
REVENUE	NO	YES	NO
REVENUEW	NO	NO	NO
VRDB90	NO	NO	NO

STORAGE GROUP SERVICE DEFINITIONS

AUTO MIGRATE

The AUTO MIGRATE value shows whether the DASD volumes in this Storage Group are eligible for automatic space management processing. Space management processing is a function of DFSMSHsm that manages DASD storage automatically by deletion of temporary datasets, release of unused, over-allocated space, and migration of low-activity datasets from client accessible volumes to HSM volumes. Primary space management process daily at 0300 and interval migration is hourly.

Possible values: YES, NO, INTERVAL, or PRIMARY

If AUTO MIGRATE is YES or INTERVAL, the datasets are eligible for primary space management and automatic interval migration. If AUTO MIGRATE is NO, the datasets are not eligible for automatic migration. If AUTO MIGRATE is PRIMARY the datasets are only eligible for primary space management. Dataset eligibility is based on the management class attributes.

AUTO BACKUP

The value in the AUTO BACKUP data column shows whether the volumes in this Storage Group are eligible for automatic backup processing. The automatic backup processing is a function of DFSMSHsm that makes incremental backup copies of newly created or changed datasets based on the management class attributes assigned to each dataset. Automatic backup processes daily at 0130.

Possible values: YES or NO

If AUTO BACKUP is YES, the datasets are eligible for automatic backup. If AUTO BACKUP is NO, then none of the datasets in this Storage Group will be backed up automatically.

AUTO DUMP

The value in the AUTO DUMP data column shows whether all volumes in this Storage Group can be automatically dumped using DFSMSHsm. The automatic dump function copies all the data from the DASD volume to tape volumes weekly beginning at 0015.

Possible values: YES or NO

STORAGE CLASS

Storage Classes are used to specify performance and availability levels for SMS-managed datasets. These are used to select a device to meet those goals and requirements thus allowing the dataset allocation to be made in the appropriate storage group. However at ITS, we do not differentiate between DASD. All production DASD is considered high performance. The storage class is used to assign the proper storage group. These may be specified in JCL, however if an incorrect storage class is specified the dataset would then be assigned to either an incorrect storage group or an incorrect management class, or both. For instance, the dataset could reside on a volume and not be backed up for disaster recovery purposes. For these reasons, CS **strongly encourages** clients to contact CS and have their datasets coded into the ACS Routines. This eliminates the need to code the STORCLAS parameter into JCL. The following table lists the defined Storage Classes currently in use. The table also shows what Storage Group would be assigned based on the Storage Class defined to the dataset within the ACS routines.

Table 5 Storage Class and corresponding Storage Group, Part 1

STORAGE CLASS	STORAGE GROUP
ACTSDB	ACTSDB
ACTSDBI	ACTSDBI
ACTSDBRO	ACTSDBRO
ACTSTEST	ACTSTEST
ACTSVRU	ACTSVRU
ARBS	ARBS
BASE	GENPROD
BASE1	TEMPSYSA
CORPERF2	CORPERF2
DBBASE	DB2PROD
DB2DASD	DB2PROD
DB2PROD	DB2PROD
DB2TEST	DB2TEST
DLDBBASE	DLDB90
DLDBTEST	DB2PROD
EISTEST	EISTEST
EISTIMS1	EISTIMS1
EISTIMS2	EISTIMS2
GENPROD	GENPROD
GENTEST	GENTEST
LARGE	LARGE
LARGET	LARGET
PERFORM	PERFORM
REVENUE	REVENUE
REVENUEW	REVENUEW
RMDS	RMDS
TSOBASE	GENPROD
VRDBBASE	VRDB90

MANAGEMENT CLASS

Management Classes for SMS-managed DASD datasets control dataset backup, retention, and migration services. DFSMSHsm uses these attributes to automatically provide both storage management and availability management. Management classes let you define management requirements for individual datasets, rather than defining the requirements for entire volumes. Management classes may be specified in JCL, but attributes are subject to change and possibly the class may be deleted. For this reason, CS **strongly encourages** clients to contact CS and have their datasets coded in the ACS routines.

CS will **implicitly** assign the following management classes via the ACS routines based on the rules set forth.

Management Class Name	Rule that assigns the Management Class to the dataset
EXTRBKP	Assigned to all DASD datasets that have a low-level(last) qualifier of COB* ASM* JCL* CNTL* CLIST LINKLIB PLI PROCLIB SCRIPT SRC SOURCE TEXT FOR* JCLLIB COBOL LOAD* PGMLIB RESLIB TEXT PL1
GONEIN3	Assigned to DASD datasets that have not been defined as SMS-managed and no volume serial has been specified in the JCL.
GONEIN3A	Assigned to DASD datasets with the naming convention: TS*.SYSA.BACKUP.** TS*.SYSB.BACKUP.** TS*.SYSE.BACKUP.**
GONEIN3W	Assigned to DASD datasets with the naming convention: TS*.SYSC.BACKUP.** TS*.SYST.BACKUP.** TS*.SYSW.BACKUP.**
NOACTION	Assigned to all DASD datasets with a storage class of 'LARGE' or 'PERFORM' or with a high-level(first) qualifier of one of the following: %%%PROD %%%TEST OR a third level qualifier of one of the following: PRODOI PRODBI TESTOI TESTBI
RMDS	This management class will be assigned to all DASD datasets with the fourth qualifier of: RMDS*
STANDARD	This management class will be assigned to all DASD datasets that have been defined as SMS-managed and there has been no specification of a special class in the client's JCL or via the ACS routines.
TSOBASE	This management class will be assigned to all DASD datasets with a high-level qualifier of: %%%TSO

MANAGEMENT CLASS ATTRIBUTES

The following three tables list all attributes of the management classes for DASD datasets. The specific classes listed are assigned via the ACS Routines. Definitions and possible values follow in *Management Class Attribute Definitions*.

Table 6 Management Class Attributes, Part 1

ATTRIBUTES	MANAGEMENT CLASS NAME				
	EXTRBKP	GONEIN3	GONEIN3A	GONEIN3W	NOACTION
Expire After Days Non-Usage	NOLIMIT	3	4	4	NOLIMIT
Expire After Date/Days	NOLIMIT	NOLIMIT	NOLIMIT	NOLIMIT	NOLIMIT
Retention Limit	NOLIMIT	0	0	0	NOLIMIT
Partial Release	COND IMMED	COND IMMED	NO	NO	NO
Primary Days Non-Usage	15	0	--	--	--
Level 1 Days Non-Usage	45	0	--	--	--
Command/Auto Migrate	BOTH	COMMAND	NONE	NONE	NONE
# GDG Elements on Primary	1	1	1	1	1
Rolled Off GDS Action	EXPIRE	EXPIRE	EXPIRE	EXPIRE	EXPIRE
Backup Frequency	0	--	--	--	--
# Backup Versions (DS Exists)	5	--	--	--	--
# Backup Versions (DS Deleted)	2	--	--	--	--
Retain Days Only Backup	60	--	--	--	--
Retain Days Extra Backup	30	--	--	--	--
Admin/User CMD Backup	BOTH	NONE	NONE	NONE	NONE
Auto Backup	YES	NO	NO	NO	NO
Backup Copy Tech	STANDARD	STANDARD	STANDARD	STANDARD	STANDARD

Table 7 Management Class Attributes, Part 2

ATTRIBUTES	MANAGEMENT CLASS NAME		
	RMDS	TSOBASE	STANDARD
Expire After Days Non-Usage	NOLIMIT	NOLIMIT	NOLIMIT
Expire After Date/Days	NOLIMIT	NOLIMIT	NOLIMIT
Retention Limit	NOLIMIT	NOLIMIT	0
Partial Release	YES IMMED	COND IMMED	COND IMMED
Primary Days Non-Usage	0	10	15
Level 1 Days Non-Usage	0	15	20
Command/Auto Migrate	BOTH	BOTH	BOTH
# GDG Elements on Primary	--	1	1
Rolled Off GDS Action	--	EXPIRE	EXPIRE
Backup Frequency	1	0	0
Backup Versions (DS Exists)	2	5	2
# Backup Versions (DS Deleted)	1	2	2
Retain Days Only Backup	5	60	60
Retain Days Extra Backup	5	30	30
Admin/User CMD Backup	BOTH	BOTH	BOTH
Auto Backup	YES	YES	YES
Backup Copy Tech	STANDARD	STANDARD	STANDARD

MANAGEMENT CLASS ATTRIBUTE DEFINITIONS

EXPIRE NON-USAGE	<p>The value in the EXPIRE NON-USAGE data column shows how many days an unaccessed dataset can exist before expiring. Datasets become eligible for expiration when the number of days since last access reaches the value in this column.</p> <p>Possible values: 1 to 9999 or NOLIMIT</p>
EXPIRE DATE/DAYS	<p>The value in the EXPIRE DATE/DAYS data column shows the expiration date or number of days until the datasets expire, beginning with the creation date.</p> <p>Possible values: yyy/mm/dd, 0 to 9999, or NO LIMIT</p> <p>If both the EXPIRE AFTER DAYS NON-USAGE and EXPIRE AFTER DATE/DAYS fields have a value of NOLIMIT, the datasets never expire. If either has a value of NOLIMIT and the other field specifies an expiration date or the number of days until expiration, the datasets expire at the specified time. If both fields specify expiration dates or the number of days until expiration, the datasets or objects expire on the later date.</p>
RET LIMIT	<p>The RET LIMIT value shows whether the DFSMSHsm will use the RETPD or EXPDT that a client or Data Class specifies for a dataset. If the value is 0, SMS will not use the specified RETPD or EXPDT. A value of 1 to 9999 represents a number of days. This number overrides RETPD or EXPDT if the specified RETPD or EXPDT exceeds it. Otherwise, SMS uses the specified RETPD or EXPDT, ignoring the number. A value of NOLIMIT allows an unlimited RETPD or EXPDT.</p> <p>Possible values: 0 to 9999 or NOLIMIT</p>

PARTIAL RELEASE

The value in the PARTIAL RELEASE data column shows whether or not the datasets can have unused space automatically released. Partial Release applies only to VSAM dataset in extended format, or non-VSAM datasets.

Possible values:

YES	Release unused space automatically during the Space Management cycle.
CONDITIONAL	Unused space can be released automatically only if a secondary allocation exists for the dataset.
YES IMMEDIATE	Release unused space when a dataset is closed or during the Space Management cycle, whichever comes first.
COND IMMEDIATE	Unused space for datasets with secondary allocation is released either when a dataset is closed or during the Space Management cycle, whichever comes first.
NO	Do not release unused space.

Note that for Extended VSAM datasets, CONDITIONAL or COND IMMEDIATE is treated the same as YES or YES IMMEDIATE.

PRIMARY DAYS

The value in the PRIMARY DAYS data column represents the minimum number of days that must elapse since last access before a dataset is eligible for migration. A value of 0 means that datasets are immediately eligible.

Possible values: 0 to 9999

LEVEL 1 DAYS

The LEVEL 1 DAYS column value indicates whether datasets can migrate to Level 1 storage and how long they can remain there.

Possible values:

0	No migration to Level 1. Datasets migrate directly from primary storage to Level 2.
1 to 9999	The total number of consecutive days that datasets must remain unaccessed before becoming eligible to migrate from Level 1 to Level 2.
NOLIMIT	Datasets cannot migrate to Level 2 automatically, but they can do so by command, or they can remain on Level 1 for an unlimited period.

CMD/AUTO MIGRATE

The CMD/AUTO MIGRATE column shows whether datasets can migrate between storage levels. The field also shows how migration, if allowed, can be initiated.

Possible values:

BOTH	Datasets can migrate either automatically or by command
------	---

	COMMAND	Datasets can migrate by command only.
	NONE	Datasets cannot migrate between storage levels.
# GDG ON PRIMARY		<p>The # GDG ON PRIMARY column shows how many of the newest generations of a GDG must always occupy primary storage. Any generations older than this set of newest generations are eligible for early migration. They will be chosen for migration in preference to non-generation datasets.</p> <p>If the field shows 0, all generations of the GDG are given priority for early migration.</p> <p>Possible values: 0 to 255</p>
ROLLED-OFF GDS ACTION		<p>The value in the ROLLED-OFF GDS ACTION data column indicates whether the GDS will expire or migrate after it has been removed from the GDG.</p> <p>Possible values: MIGRATE or EXPIRE</p>
BACKUP FREQUENCY		<p>A value of 1 to 9999 in the BACKUP FREQUENCY data column represents the minimum number of days between backups for datasets. A new backup of a dataset can be made after this period of days only if the dataset is changed during that period.</p> <p>A BACKUP FREQUENCY of 0 means that the dataset, if changed, will be backed up each time DFSMSHsm processes the volume that contains them.</p> <p>Possible values: 0 to 9999</p>
# BACKUPS (DS EXISTS)		<p>The # BACKUPS (DS EXISTS) value shows the maximum number of backups versions to retain for a dataset.</p> <p>Only the most recent automatic backups can be kept. Each backup of a given dataset will contain a different version of the dataset.</p> <p>Possible values: 1 to 13</p>
# BACKUPS (DS DELETED)		<p>The # BACKUPS (DS DELETED) value shows whether automatic backups of a dataset will be kept after the dataset is deleted. A value of 0 means that no such backups will be kept. A value of 1 or higher represents the maximum number that can be kept.</p> <p>Each automatic backup of a deleted dataset contains a different version of the dataset. Only the most recent backups will be kept.</p> <p>Possible values: 0 to 13</p>
RETAIN DAYS ONLY BACKUP		<p>The value in the RETAIN DAYS ONLY BACKUP data column shows how long the most recent backup copy of a dataset will be kept after the dataset is deleted. A numeric value represents a specific number of days; NOLIMIT means unlimited retention.</p> <p>Possible values: 1 to 9999 or NOLIMIT</p>
RETAIN DAYS EXTRA		

BACKUPS

The value in the RETAIN DAYS EXTRA BACKUPS column shows how long to keep backups of a dataset that pre-date the most recent backup. Each of these older backups will be kept for the period specified, regardless of whether the dataset exists or has been deleted. A numeric value represents a specific number of days; NOLIMIT means unlimited retention.

Possible values: 1 to 9999 or NOLIMIT

ADM/USER BACKUP

The value in the ADM/USER BACKUP column shows who is authorized to perform command backups against the datasets. If only the Storage Administrator can perform command backups, the value will be ADMIN. BOTH means that clients as well as the Storage Administrator can perform command backups. NONE means that neither clients nor the Storage Administrator can perform command backups.

Possible values: ADMIN, BOTH or NONE

AUTO BACKUP

The value in the AUTO BACKUP column shows whether automatic backup is allowed for datasets.

Possible values: YES or NO

EXAMPLES OF DATASET NAMES

This section gives examples of dataset names, how those datasets are stored and managed by SMS, and the page number in the preceding documentation that discusses each allocation.

DB2 DATASETSTEST**ADMDB51.DSNDBD.GPSD000.ILBS0A00.I0001.B001**

Data Class:	DEFAULT	(see page 12)
Management Class:	NOACTION	(see page 34)
Storage Class:	DB2TEST	(see page 32)
Storage Group:	DB2TEST	(see page 26)

PRODUCTION**CORDB05.DSNDBD.CORDB01.ASII8AA2.I0001.A001**

Data Class:	DEFAULT	(see page 12)
Management Class:	NOACTION	(see page 34)
Storage Class:	DBBASE	(see page 32)
Storage Group:	DB2PROD	(see page 26)

TSO DATASETS

TSO datasets are not differentiated between Production and Test.

ADMTSO . TEAM01 . JCL

Data Class:
Management Class: EXTRBK (see page 34)
Storage Class: TSOBASE (see page 33)
Storage Group: GENPROD (see page 27)

TRETSO . TEAM46 . PARMLIB

Data Class:
Management Class: TSOBASE (see page 35)
Storage Class: TSOBASE (see page 33)
Storage Group: GENPROD (see page 26)

GENERAL DATASETS

PRODUCTION

CCD . CCA . PRODBS . CCA103 - 0 . ALAMANCE . SPRING98

Data Class: DATAF (see page 11)
Management Class: NOELSE (see page 35)
Storage Class: BASE (see page 32)
Storage Group: GENPROD (see page 26)

TEST

DHR . CAP . TESTBS . PDCLAIMS . AP110 . MAR99

Data Class:	DATAF	(see page 11)
Management Class:	STANDARD	(see page 35)
Storage Class:	GENTEST	(see page 32)
Storage Group:	GENTEST	(see page 27)

HIERARCHICAL STORAGE MANAGER (HSM)

Hierarchical Storage Manager provides space management, backup, and recovery functions to manage datasets automatically on a variety of storage devices. It thus reduces manual intervention and optimizes the use of primary storage space, such as G93 volumes. It does so by moving (migrating) and compacting "aged" datasets and then automatically recalling them to primary volumes when they are referenced by a batch job or needed during a TSO session.

After HSM has migrated datasets, the System Catalog indicates that they are cataloged by the volume MIGRAT. Although there is physically no such volume mounted on the system, HSM operates as if there were. For example, if you reference migrated datasets via a standard TSO or batch job command, HSM will recall the datasets for you and then will process your TSO command or proceed with the batch job.

HSM COMMANDS

HLIST

The HLIST command lists information about migrated and/or backed up datasets using the MCDS or the BCDS. If you do not specify which type of information is desired, HSM will list migrated datasets.

The following will list information about the specified dataset from both the MCDS and the BCDS:

HLIST BOTH DATASETNAME(*dataset name*)

The following will list information about the specified backed up dataset from the BCDS to the terminal:

HLIST BCDS DATASETNAME(*dataset name*)

The following will list information about datasets beginning with the HLQ within the MCDS to your terminal:

HLIST LEVEL(*high level qualifier*)

The following will list information about datasets beginning with the HLQ within the MCDS to a permanent dataset:

HLIST LEVEL(*high level qualifier*) ODS(*dataset name*)

The results of the HLIST command default to the terminal unless the ODS is specified.

HRECOVER

The command HRECOVER is used to recover a backup version of one or more datasets.

The following will recover a dataset from the most recent backup generation:

HRECOVER *dataset name*

The following will recover a dataset from the X backup volume generation:

HRECOVER *dataset name* GEN(*xxx*)

In this command, *xxx* represents the number of the GDG that you are recovering.

HRECALL

The following will recall a dataset from a migration volume to a primary disk volume without tying up the TSO session.

HRECALL *dataset name* NOWAIT

HDELETE

HDELETE deletes one or more migrated datasets from a migration volume. HSM deletes the dataset without recalling it to a primary volume. When HSM deletes the dataset it maintains any backup versions of the dataset and the information in the BCDS.

HDELETE cannot be used to delete datasets from primary volumes or backup volumes.

The following deletes a migrated dataset from the MCDS without waiting for the command to complete:

HDELETE *dataset name* NOWAIT

TAPE MANAGEMENT

Section to be included in the future.