

z/OS



DFSMS: Using the Volume Mount Analyzer

z/OS



DFSMS: Using the Volume Mount Analyzer

Note

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 181.

Second Edition, September 2004

This edition applies to Version 1 Release 6 of z/OS[®] (5694-A01), Version 1 Release 6 of z/OS.e (5655-G52), and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC26-7342-00.

IBM[®] welcomes your comments. A form for readers’ comments may be provided at the back of this publication, or you may address your comments to the following address:

- International Business Machines Corporation
- Department 55JA, Mail Station P384
- 2455 South Road
- Poughkeepsie, NY 12601-5400
- United States of America
-
- FAX (United States & Canada): 1+845+432-9405
- FAX (Other Countries):
- Your International Access Code +1+845+432-9405
-
- IBMLink[™] (United States customers only): IBMUSM10(MHVRFCFS)
- Internet e-mail: mhvrdfs@us.ibm.com
- World Wide Web: <http://www.ibm.com/servers/eserver/zseries/zos/webqs.html>

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:

- Title and order number of this document
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1987, 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	vii
--------------------------	------------

Tables	ix
-------------------------	-----------

About This Book.	xi
-----------------------------------	-----------

Required Product Knowledge	xi
--------------------------------------	----

Referenced Publications	xii
-----------------------------------	-----

Summary of Changes	xiii
-------------------------------------	-------------

Summary of Changes for SC26-7413-01 z/OS	
--	--

Version 1 Release 6	xiii
-------------------------------	------

Chapter 1. Preparing to Use the Volume Mount Analyzer **1**

Understanding the Volume Mount Analyzer	1
---	---

Using the Volume Mount Analyzer for Tape Mount Management	1
---	---

What Information Does the Volume Mount Analyzer Give You?	2
---	---

How Do You Use This Information for Tape Mount Management?	2
--	---

How Do You Manage Your Tape Data with Tape Mount Management?	2
--	---

What Are the Benefits of Using Tape Mount Management?	3
---	---

Performing a Volume Mount Analyzer Study	3
--	---

Collecting SMF Records	4
----------------------------------	---

Understanding the Limitations of the Volume Mount Analyzer.	5
---	---

Shared Systems: Same Addresses for Tape Subsystems.	5
---	---

Shared Systems: Different Addresses for Tape Subsystems.	6
--	---

Avoiding Confusion with Shared Tape Subsystems	7
--	---

Knowing Notational Conventions	7
--	---

Chapter 2. Using GFTAXTR to Extract Data from SMF Records. **9**

Understanding the GFTAXTR Program	9
---	---

GFTAXTR Objectives	9
------------------------------	---

GFTAXTR Processing	9
------------------------------	---

Understanding the Input to GFTAXTR	10
--	----

Required SMF Records	10
--------------------------------	----

SMF Record Types	10
----------------------------	----

Submitting the GFTAXTR JCL	11
--------------------------------------	----

Submitting the GFTAXTR JCL for Concatenated Input	14
---	----

Using GFTAXTR Keywords	14
----------------------------------	----

Keyword Summary and Syntax.	15
-------------------------------------	----

Keyword Rules	15
-------------------------	----

ACCTFLD Keyword	15
---------------------------	----

DATE Keyword	16
------------------------	----

LINES Keyword	17
-------------------------	----

TIME Keyword	17
------------------------	----

Interpreting the GFTAXTR Output Report	18
--	----

GFTAXTR Output Report—Input Phase	18
---	----

GFTAXTR Output Report—Messages.	19
---	----

GFTAXTR Output Report—Internal Record Generation Phase	19
--	----

GFTAXTR Output Report—Internal SMF	
------------------------------------	--

Job/Step Reconstruction Phase	21
---	----

GFTAXTR Output Report—Data Set Statistics Phase	23
---	----

GFTAXTR Output Report—NON-HSM Tape Data Set Usage Analysis	24
--	----

GFTAXTR Output Report—Final Record Output Phase	26
---	----

Examining the GFTAXTR Output	27
--	----

Troubleshooting when Obtaining the GFTAXTR Output	28
---	----

Chapter 3. Understanding Tape—Some Concepts. **29**

Single/Multi Volume/File Data Sets	29
--	----

Single Volume and Single File	29
---	----

Single Volume and Multifile	29
---------------------------------------	----

Multivolume and Single File.	29
--------------------------------------	----

Multivolume and Multifile	30
-------------------------------------	----

Statistical Mounts	30
------------------------------	----

Example of Calculating Statistical Mounts	30
---	----

Example for Multifile Data Sets.	31
--	----

Example for Multiple Versions of a Data Set	31
---	----

Statistical Volumes	31
-------------------------------	----

Unique Volume Serial and Statistical Volumes	32
--	----

Zero Statistical Volumes	33
------------------------------------	----

Example of Calculating Statistical Volumes	33
--	----

Example for Multiple Versions of a Data Set	33
---	----

Example for Multifile and Multiple Version Data Sets	34
--	----

Open Counts	34
-----------------------	----

Tape Cartridge Usage	35
--------------------------------	----

Volumes Without IDRC	35
--------------------------------	----

Cartridges With IDRC	35
--------------------------------	----

Data Set Usage Types	36
--------------------------------	----

Order for Assigning Usage Types	36
---	----

HSM	36
---------------	----

Temporary.	37
--------------------	----

Backup.	37
-----------------	----

Single	37
------------------	----

BCOPY.	37
----------------	----

Active	38
------------------	----

Categories of Tape Mount Management Data Sets	38
---	----

Order for Assigning Tape Mount Management	
---	--

Data Set Categories.	38
------------------------------	----

Tape Mount Management Data Set Categories in	
--	--

the Estimate Report.	39
------------------------------	----

Categorization of Data Sets	40
---------------------------------------	----

Chapter 4. Analyzing Your Tape Usage with GFTAVMA 41

Understanding the GFTAVMA Program	41
Requirements for Using GFTAVMA	41
GFTAVMA Output Reports	41
Understanding GFTAVMA Processing Phases	42
Preparing to Run GFTAVMA	43
Filtering Your Input with GFTAVMA Keywords	44
Filter Types	44
Categories of Include/Exclude Filters	44
Example of Primary and Secondary Filters	45
Include/Exclude Filter Syntax	45
Global Data Set Name Characters in Filters	46
Examples of Global Data Set Name Characters	46
Examples of Include/Exclude Filters	46
Example of Generation Data Set Filters	47
Example of Multiple Filters	47
Filter Cautions	48
Submitting the GFTAVMA JCL	48
Troubleshooting When Obtaining GFTAVMA Output	51
Interpreting GFTAVMA Summary Reports	51
Input Analysis Phase—GFTASRT1	52
Volume Analysis Phase—GFTASRT2	53
Data Set Analysis Phase—GFTASRT3	57

Chapter 5. Using GFTAVMA Keywords to Filter Your Input 63

Keyword Rules	63
Summary of Keywords That Filter Your Input	63
ACCOUNT Keyword	64
Syntax	64
Default	65
Relationship to Other Keywords	65
ACCOUNT Examples	65
Using ACCOUNT as an ACS Filter	65
DATASET Keyword	65
Syntax	65
Default	65
DATASET Examples	66
Using DATASET as an ACS Filter	66
DATE Keyword	66
Syntax	66
Default	67
DATE Examples	67
EXPDT Keyword	67
Syntax	67
Default	68
EXPDT Examples	68
Using EXPDT as an ACS Filter	68
FILE Keyword	68
Syntax	68
Default	69
Relationship to Other Keywords	69
FILE Examples	69
JOBNAME Keyword	69
Syntax	69
Default	70
JOBNAME Example	70
Using JOBNAME as an ACS Filter	70
MAXSIZE and MINSIZE Keywords	70

Syntax	70
Default	70
Caution	70
Relationship to Other Keywords	71
MAXSIZE and MINSIZE Examples	71
MOUNT Keyword	72
Syntax	72
Default	72
Relationship to Other Keywords	72
MOUNT Examples	72
PROGRAM Keyword	73
Syntax	73
Default	73
PROGRAM Examples	73
Using PROGRAM as an ACS Filter	73
SYSTEMID Keyword	74
Syntax	74
Default	74
SYSTEMID Example	74
TIME Keyword	74
Syntax	74
Default	75
TIME Examples	75
UNIT Keyword	75
Syntax	75
Default	76
Relationship to Other Keywords	76
UNIT Examples	76
UNITADDR Keyword	76
Syntax	77
Default	77
Relationship to Other Keywords	77
Caution	77
UNITADDR Examples	77
USAGE Keyword	78
Syntax	78
Default	79
VOLUME Keyword	79
Syntax	79
Default	79
Relationship to Other Keywords	79
Restrictions	80
VOLUME Example	80

Chapter 6. Producing Detailed Analysis Reports 81

Understanding Detailed Analysis Reports	81
Production Phases of Detailed Reports	81
Types of Detailed Reports	81
Interpreting the Actual Reports	82
The Data Set Report	82
The Maximum Gigabyte Report	88
The Top Report	92
The Usage Report	95
The Volume Report	101
Interpreting the Simulation Reports	106
The Estimate Report	106
The Improved Data Recording Capability Report	120
The Management Class Report	125

Chapter 7. Using GFTAVMA Keywords to Tailor Your Output 131

Keyword Rules	131	Syntax	143
Summary of Keywords That Tailor Your Output Reports	131	GFTAVMA and Compacted Data Sets	143
BLKSIZE Keyword	132	Default	143
Syntax	133	Relationship to Other Keywords	143
Default	133	Reports Affected	144
BLKSIZE Example	133	LARGE Keyword	144
Reports Affected	133	Syntax	144
BPROGRAM Keyword	133	Default	144
Syntax	134	Relationship to Other Keywords	144
Default	134	Reports Affected	145
Relationship to Other Keywords	134	LEVELOAGE Keyword	145
Reports Affected	134	Syntax	145
CHART Keyword	134	Default	145
Syntax	134	Relationship to Other Keywords	145
Additional JCL Needed for CHART	135	Reports Affected	146
Default	135	LEVELIAGE Keyword	146
CHART Example	135	Syntax	146
Reports Affected	135	Default	146
CLASSIFY Keyword	135	Relationship to Other Keywords	146
Syntax	135	Reports Affected	147
Default	136	LINES Keyword	147
Relationship to Other Keywords	136	Syntax	147
Restrictions	136	Default	147
Reports Affected	136	Reports Affected	147
COMPACTION Keyword	136	L1FSPACE Keyword	147
Syntax	137	Syntax	147
Default	137	Default	148
Restrictions	137	Reports Affected	148
COMPACTION Example	137	REPORT Keyword	148
Reports Affected	138	Syntax	148
CUTOFF Keyword	138	Parameters for the Top Report	149
Syntax	138	Default	149
Default	138	Relationship to Other Keywords	149
CUTOFF Example	138	REPORT Examples	149
Reports Affected	138	Reports Affected	150
DASDDEV Keyword	138	RPROGRAM Keyword	150
Syntax	139	Syntax	150
Default	139	Default	151
DASDDEV Example	139	Reports Affected	151
Reports Affected	139	SPLIT Keyword	151
DSORT Keyword	139	Syntax	151
Syntax	139	The Importance of SPLIT	151
Default	140	Additional JCL Needed for SPLIT	152
Reports Affected	140	Default	152
FORCECAP/NOFORCECAP Keywords	140	Reports Affected	152
Syntax	140	TAPEDEV Keyword	152
The Importance of FORCECAP	141	Syntax	152
Default	141	Default	153
Reports Affected	141	Relationship to Other Keywords	153
FSPACE Keyword	141	Reports Affected	153
Syntax	141	TAPELEN Keyword	153
Default	142	Syntax	153
Reports Affected	142	Default	154
GDG Keyword	142	Relationship to Other Keywords	154
Syntax	142	Reports Affected	154
Default	142	VSORT Keyword	154
Reports Affected	142	Syntax	154
IDRCFACTOR Keyword	143	Default	154
		Reports Affected	155

**Appendix A. A Quick Reference to
Volume Mount Analyzer Keywords . . . 157**

Abbreviations and Defaults for GFTAXTR
Keywords 157
Abbreviations and Defaults for Keywords That
Filter Your Input 157
Abbreviations and Defaults for Keywords That
Tailor Output Reports 158

**Appendix B. Volume Mount Analyzer
Keyword Syntax 159**

Summary of GFTAXTR Keyword Syntax 159
Summary of GFTAVMA Keyword Syntax 159

Appendix C. Example of Chart Output 161

**Appendix D. Volume Mount Analyzer
Messages 167**

Notices 181
Trademarks 182

Glossary 183

Index 189

Figures

1. ACCTFLD Keyword Syntax	15	36. The Data Set Report Legend	84
2. DATE Keyword Syntax	16	37. Partial Data Set Report.	87
3. LINES Keyword Syntax	17	38. The GBMAX Report.	91
4. TIME Keyword Syntax.	17	39. The Top Report	94
5. Example GFTAXTR Report—Input Phase	18	40. The Usage Report—Tape Allocations, System	
6. Example GFTAXTR Report—Internal Record		Statistics	97
Generation Phase (Messages).	19	41. The Usage Report—Tape Allocations by Hour	98
7. Example GFTAXTR Report—Internal Record		42. The Usage Report—Tape Allocations by Hour	
Generation Phase	20	(continued).	99
8. Example GFTAXTR Report—Internal SMF		43. The Usage Report—Tape Mounts	100
JOB/STEP Reconstruction Phase	22	44. The Usage Report—Tape Mounts (continued)	101
9. Example GFTAXTR Report—Data Set Statistics		45. The Volume Report Legend	103
Phase	23	46. Partial Volume Report	105
10. Example GFTAXTR Report—Non-HSM Tape		47. Estimate Report—Statistical Mount Savings	
Data Set Usage Analysis	25	Breakdown	108
11. Example GFTAXTR Report—Final Record		48. Estimate Report—Statistical Volume Savings	
Output Phase	26	Breakdown	111
12. Count of unique volsers for GFTASRT2	32	49. Estimate Report—DASD Buffer Cost	
13. Count of statistical volumes for GFTASRT3	33	Breakdown	113
14. Order of Assigning Tape Mount Management		50. The Estimate Report Legend	116
Categories	40	51. The Estimate Report—Savings/Cost Summary	118
15. The Three Phases of GFTAVMA Processing	42	52. The IDRC Report Legend	123
16. GFTASRT1 Summary Report—Input Analysis		53. The IDRC Report	124
Phase	53	54. Partial Management Class Report.	127
17. GFTASRT2 Summary Report—Tape Volser		55. BLKSIZE Keyword Syntax	133
Processing Phase (The # character in Volume		56. BPROGRAM Keyword Syntax	134
Mount Analyzer reports represents the word		57. CHART Keyword Syntax	135
"number").	54	58. CLASSIFY Keyword Syntax.	136
18. GFTASRT2 Summary Report—Tape Volume		59. COMPACTION Keyword Syntax	137
Size Ranges	56	60. CUTOFF Keyword Syntax	138
19. GFTASRT3 Summary Report—Tape Data Set		61. DASDDEV Keyword Syntax	139
Processing Phase.	58	62. DSORT Keyword Syntax.	139
20. GFTASRT3 Summary Report—Tape Data Set		63. FORCECAP/NOFORCECAP Keyword Syntax	140
Analysis	60	64. FSPACE Keyword Syntax	142
21. ACCOUNT Keyword Syntax.	64	65. GDG Keyword Syntax	142
22. DATASET Keyword Syntax	65	66. IDRCFACTOR Keyword Syntax	143
23. DATE Keyword Syntax	66	67. LARGE Keyword Syntax.	144
24. EXPDT Keyword Syntax	67	68. LEVEL0AGE Keyword Syntax	145
25. FILE Keyword Syntax	68	69. LEVEL1AGE Keyword Syntax	146
26. JOBNAME Keyword Syntax	69	70. LINES Keyword Syntax	147
27. MAXSIZE and MINSIZE Keyword Syntax	70	71. L1FSPACE Keyword Syntax.	147
28. MOUNT Keyword Syntax.	72	72. REPORT Keyword Syntax	148
29. PROGRAM Keyword Syntax.	73	73. RPROGRAM Keyword Syntax	150
30. SYSTEMID Keyword Syntax	74	74. SPLIT Keyword Syntax	151
31. TIME Keyword Syntax.	75	75. TAPEDEV Keyword Syntax.	152
32. UNIT Keyword Syntax.	76	76. TAPELEN Keyword Syntax	153
33. UNITADDR Keyword Syntax	77	77. VSORT Keyword Syntax	154
34. USAGE Keyword Syntax	78	78. Example Chart Output	163
35. VOLUME Keyword Syntax	79		

Tables

1. SMF Input Records	11	8. Quick Reference for Keywords that Filter Your Input	157
2. Keyword Summary for GFTAXTR	15	9. Quick Reference for Keywords that Tailor Your Output Reports	158
3. Capacities of Tape Volumes without IDRC Compaction	35	10. Keyword Summary for GFTAXTR	159
4. Capacities of Tape Cartridges with an IDRC Compaction Factor of 3	35	11. Keyword Summary for GFTAVMA	159
5. Keywords that Filter Your Input	63	12. Severity Levels for the Volume Mount Analyzer Messages	167
6. Keywords that Tailor Your Output Reports	131		
7. Quick Reference to GFTAXTR Keywords	157		

About This Book

This book is intended for storage administrators who run volume mount analyzer programs and use the generated reports to improve the efficiency of tape storage.

This book describes how to use the volume mount analyzer to determine the value of automating tape data management using z/OS™. The volume mount analyzer produces reports that show the mounts and tape usage at your installation. You can use these reports to decide whether your installation could benefit from tape mount management.

Specifically, this book provides information to help you:

- Prepare to run the volume mount analyzer
- Extract data from your SMF records to provide input to the volume mount analyzer
- Understand tape concepts
- Run the volume mount analyzer to study your tape usage and produce summary reports
- Use keywords to filter your input
- Produce optional detailed volume mount analyzer reports
- Use keywords to tailor your output reports
- Understand the volume mount analyzer messages

| Before using this book, see *z/OS DFSMS: Implementing System-Managed Storage* for
| an explanation of how to implement tape mount management using information
| from the volume mount analyzer reports.

Required Product Knowledge

To use this book effectively, you should be familiar with:

- Applications that use tape at your installation
- DFSMS™
- Method of allocation in MVS
- Job control language (JCL)
- System management facilities (SMF) records
- Tape and DASD hardware
- Tape mount management

| You should also familiarize yourself with related information presented in the
| following publications:

- *z/OS DFSMS: Implementing System-Managed Storage*, SC26-7407
- *z/OS MVS JCL Reference*, SA22-7597
- *z/OS MVS JCL User's Guide*, SA22-7598
- *z/OS MVS System Management Facilities (SMF)*, SA22-7630

| For more specific information on the Storage Management Subsystem, see *z/OS*
| *DFSMSdfp Storage Administration Reference*, SC26-7402.

Referenced Publications

Within this book, references are made to the following publications:

- *z/OS DFSMS: Implementing System-Managed Storage*, SC26-7407
- *z/OS DFSMS: Using Magnetic Tapes*, SC26-7412
- *z/OS DFSMSHsm Implementation and Customization Guide*, SC35-0418
- *z/OS MVS JCL Reference*, SA22-7597
- *z/OS MVS JCL User's Guide*, SA22-7598
- *z/OS MVS System Management Facilities (SMF)*, SA22-7630

Summary of Changes

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

Summary of Changes for SC26-7413-01 z/OS Version 1 Release 6

| This document contains information previously presented in *z/OS DFSMS: Using*
| *the Volume Mount Analyzer*, SC26-7413-00.

| This document does not differ in technical content from the previous release. It
| was updated to correct a formatting error in the previous edition. References to
| other publications and information on submitting readers' comments have also
| been updated.

Chapter 1. Preparing to Use the Volume Mount Analyzer

This chapter provides an overview of the *volume mount analyzer* and explains how you can prepare to use it.

Understanding the Volume Mount Analyzer

The volume mount analyzer is a program that helps you analyze your current tape environment. You use the volume mount analyzer to study tape mount activity, monitor tape media use, and implement *tape mount management* at your installation. The volume mount analyzer produces reports that you can use to tailor *data classes*, *management classes*, and *automatic class selection (ACS) routine* filters depending on your specific tape usage profile. Using tape mount management, you can maximize the use of your tape media and reduce your tape mounts.

The volume mount analyzer is composed of two programs: GFTAXTR and GFTAVMA.

GFTAXTR

Processes *system management facilities (SMF)* records and puts them in a format that the GFTAVMA program can use. GFTAXTR produces an output data set, which is the input to GFTAVMA, and a report of its processing phases.

GFTAVMA

Analyzes the tape usage and tape mounts at your installation using the output data set from GFTAXTR. GFTAVMA produces summary reports for every run and optional detailed reports if you request them.

See Chapter 2, “Using GFTAXTR to Extract Data from SMF Records,” on page 9 and Chapter 4, “Analyzing Your Tape Usage with GFTAVMA,” on page 41 for more information about these two programs.

By analyzing SMF records of your installation, the volume mount analyzer can help you determine:

- the space each tape data set uses on a *tape volume*
- the unused space on a tape volume
- the number of tape mounts you would save by implementing tape mount management
- the tape mount management DASD buffer space you would need to hold the data sets until they are transferred to tape storage

See *z/OS MVS System Management Facilities (SMF)* for more information on SMF records.

Using the Volume Mount Analyzer for Tape Mount Management

The volume mount analyzer helps you analyze the costs and savings of automating your tape data management using DFSMS. It also helps you choose data sets for tape mount management. Tape mount management is a method of reducing tape mounts by:

- Allowing the system to manage the placement of data

- Taking advantage of hardware and software compaction
- Taking advantage of new tape technology
- Automatically filling each tape volume to capacity

Using DFSMS and tape mount management can help you reduce the number of both tape mounts and tape volumes that your installation requires. The volume mount analyzer reviews your tape mounts and creates reports that provide you with information you need to effectively implement the tape mount management methodology recommended by IBM®.

See *z/OS DFSMS: Implementing System-Managed Storage* and *MVS/ESA SML: Managing Storage Groups* for more information on tape mount management.

What Information Does the Volume Mount Analyzer Give You?

The volume mount analyzer produces both summary and detailed reports of your tape activity for a selected time period. You can use these reports, along with an estimate of the average cost of a tape mount, to determine the value of managing tape data with DFSMS or to write ACS routines that implement tape mount management most effectively.

The summary reports are generated with every run, but you must request the detailed reports using GFTAVMA keywords. GFTAVMA generates two types of optional detailed reports:

- Reports that show the actual tape usage at your installation
- Reports that simulate what the tape usage at your installation might be if you used tape mount management

See “Interpreting the Actual Reports” on page 82 and “Interpreting the Simulation Reports” on page 106 for more information.

All volume mount analyzer messages appear in the reports. See Appendix D, “Volume Mount Analyzer Messages,” on page 167 for more information.

How Do You Use This Information for Tape Mount Management?

Use the volume mount analyzer reports to identify the data sets that you can manage with tape mount management and to develop the ACS filters that you can use to capture and redirect these data sets to the *Storage Management Subsystem (SMS)*. The reports also show the potential costs and savings of tape mount management.

See *z/OS DFSMS: Implementing System-Managed Storage* for details on how to use the volume mount analyzer reports for tape mount management.

How Do You Manage Your Tape Data with Tape Mount Management?

To use your tape media effectively and achieve optimum benefits from advancements in compaction and tape cartridges, you must write full cartridges of data automatically. To do so, you can use DFSMS to stack small data sets on a cartridge without dependencies on manual file markers or JCL changes.

Tape mount management recommends that you use *DFSMSHsm™* to do *interval migration* to SMS *storage groups*. You can use ACS routines to redirect your tape data sets to a tape mount management DASD buffer storage group. DFSMSHsm

scans this buffer on a regular basis and migrates the data sets to *migration level 1* DASD or *migration level 2* tape as soon as possible, based on the management class and storage group specifications.

By using SMS to direct data sets from tape to DASD, you do not have to change your JCL to modify the UNIT= parameter.

Managing tape data automatically with DFSMS implies that:

- Random tape mounts are drastically reduced. DFSMS creates data in a primary pool buffer and puts it to tape at a later time using DFSMSHsm. A later time can be minutes or days, depending on the management class assigned to the particular data set.
- When a tape volume is mounted, it is filled with data. The system automatically determines the amount of data that can reside on a single volume based on the capability of both the installed hardware and the tape cartridge.
- If an application accesses a data set later, DFSMSHsm automatically recalls it from where it resides in the *storage hierarchy* and allocates it on primary DASD for access. No JCL changes are required, and the location of the data is transparent to the application program.

What Are the Benefits of Using Tape Mount Management?

Tape mount management allows you to efficiently fill a tape cartridge to its capacity and gain full benefit from *improved data recording capability (IDRC)* compaction, 3490E Enhanced Capability Magnetic Tape Subsystem, 36-track enhanced recording format, and *Enhanced Capacity Cartridge System Tape*. By filling your tape cartridges, you reduce your tape mounts and even the number of tape volumes you need.

With an effective tape cartridge capacity of 2.4 GB using 3490E and the Enhanced Capacity Cartridge System Tape, DFSMS can intercept all but extremely large data sets and manage them with tape mount management. By implementing tape mount management with DFSMS, you might reduce your tape mounts by 60 to 70% with little or no additional hardware required. Therefore, the resulting tape environment would be able to fully exploit integrated cartridge loaders (ICL), IDRC, and 3490E.

Tape mount management also improves job throughput because jobs are no longer queued up on tape drives. Approximately 70% of all tape data sets queued up on drives are less than 10MB. With tape mount management, these data sets reside on DASD while in use. This frees up the tape drives for other allocations.

Do not confuse the process of automating tape data management using DFSMS with tape-to-DASD conversion. DASD storage is used only as a temporary buffer. Little of the data currently on tape really belongs on DASD. Most often, *inactive data* is put on tape. You need to decide the best time and the most efficient way to store the data there.

Performing a Volume Mount Analyzer Study

To run the volume mount analyzer, you must have:

- *Data Facility Sort (DFSORT)* or its functional equivalent
- A z/OS operating system

In a volume mount analyzer study, you must:

1. Determine a representative time period for the study.

To get an accurate picture of tape mount activity, you should choose a time that includes peak loads as well as routine processing. A good starting point is to study a one month period, including weekends. Also, include month-end activity in the sample period. In any event, collect SMF records during a period that is representative of your tape usage.
2. Collect SMF records by enabling and creating the necessary SMF records during the chosen time period.

To perform an analysis, the volume mount analyzer uses SMF data that your installation has collected over the representative time period, which should be one month. See “Understanding the Input to GFTAXTR” on page 10 for information about the SMF records required to perform the analysis.

To run the volume mount analyzer, you must provide the required SMF records. The volume mount analyzer needs these SMF records to generate the desired output reports.
3. Run the SMF data extractor program, GFTAXTR, once to reduce the amount of input data.

You use the resulting output data set for the rest of the study unless you want to study another period of time.
4. Run the volume mount analyzer program, GFTAVMA, to generate both the summary and detailed reports.

Before you run GFTAVMA, you need to understand the data set naming conventions at your installation in order to properly categorize your data. Although the volume mount analyzer can use its built-in intelligence to categorize data, its assumptions might not be completely accurate.

To determine the potential savings and costs, you are recommended to run a simple Estimate Report (see “The Estimate Report” on page 106 for more information). Then, with subsequent GFTAVMA runs, you can isolate certain applications based on appropriate criteria, such as program name or data set name. A suggested implementation usually contains INCLUDEs of program names and, optionally, some EXCLUDEs of data set names. See “Include/Exclude Filters” on page 44 for more information.
5. Analyze the results.

At this time, construct the necessary **INCLUDE** or **EXCLUDE** filters and use the volume mount analyzer keywords. You can run GFTAVMA as many times as needed, until you are satisfied with the results.
6. Determine whether you want to implement tape mount management.

Collecting SMF Records

The system management facilities (SMF) create records of MVS system and job-related information for your installation.

To run the volume mount analyzer, you must provide SMF records for your system. When collecting SMF records, you must:

- Use SMF records from a currently supported release of z/OS.
- Use the IFASMFDP dump utility to collect the records.
- Collect the correct record types and sub-types. See “Understanding the Input to GFTAXTR” on page 10 for the required types and sub-types
- Collect the data that you want to analyze together. Once GFTAXTR has created the output file, you cannot sub-divide or merge the output with other GFTAXTR output files

- Be aware of shared systems considerations. See “Understanding the Limitations of the Volume Mount Analyzer” for more information.

For best results, you should:

- Collect SMF records for a period of at least one month, including month-end processing.
- End the collection sample just before midnight so that you have a full day of records for the last day.

See *z/OS MVS System Management Facilities (SMF)* for details on SMF records and how to collect them.

Understanding the Limitations of the Volume Mount Analyzer

GFTAXTR rebuilds tape data set and job activity from SMF records, and GFTAVMA analyzes and reports on this SMF information. The volume mount analyzer works well with standard tape usage, but it might not correctly interpret SMF data when it is merged from multiple systems.

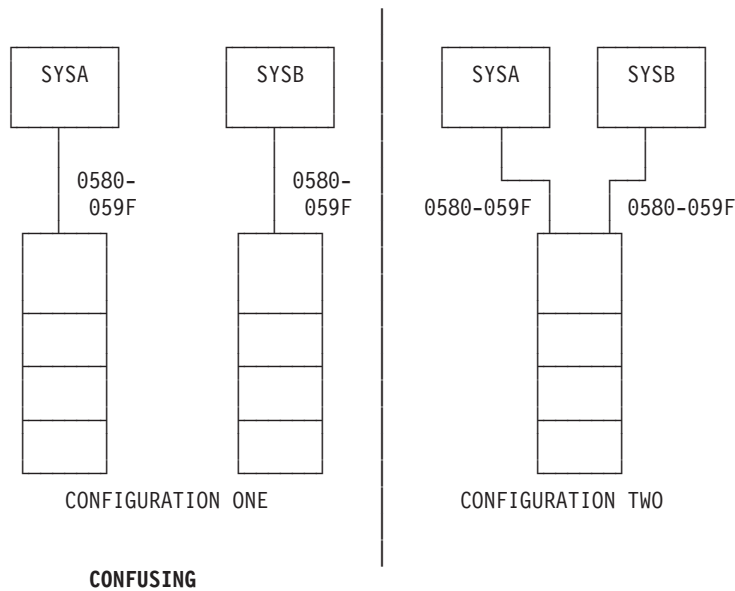
Although the volume mount analyzer can process SMF data from multiple systems, it might not understand some of the addressing configurations. Consequently, these configurations might cause misleading or incorrect results in the Usage Report. See “The Usage Report” on page 95 for more information. Therefore, you must be aware of how your tape drives are addressed when merging SMF data from multiple systems into one volume mount analyzer analysis.

The volume mount analyzer assumes that each unique device address represents one unit. This might cause problems if two systems use the same address for two different tape subsystems. Similarly, problems might occur if two systems use different addresses for the same set of shared tape subsystems.

When you use the volume mount analyzer, you must understand how *tape subsystems* are addressed and shared to avoid confusion. Look at your system configurations map to find out which types of configuration you have.

Shared Systems: Same Addresses for Tape Subsystems

When reading SMF data for systems A and B, the volume mount analyzer cannot make a distinction between Configurations One and Two.



Configuration One

SYSA and SYSB do not share tape subsystems. However, each system has tape units with the same addresses, 0580-059F. In this case, the volume mount analyzer incorrectly assumes there is one tape subsystem.

Configuration Two

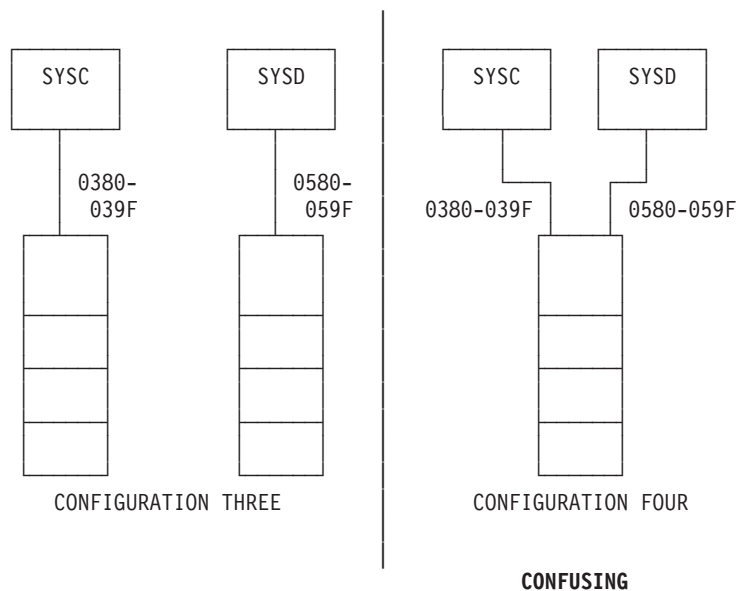
SYSA and SYSB share a tape subsystem. The volume mount analyzer correctly interprets that both systems use the same address range (0580-059F) to refer to one tape subsystem.

Therefore, merging SMF data for SYSA and SYSB in Configuration One might cause misleading or incorrect results for the Usage Report because the volume mount analyzer assumes that each unique address represents one and only one unit. See "The Usage Report" on page 95 for more information.

Note: Do not merge SMF data from multiple systems if your installation has Configuration One. Analyze these SMF samples separately.

Shared Systems: Different Addresses for Tape Subsystems

The volume mount analyzer can not distinguish between Configurations Three and Four either when you combine the SMF data from SYSC and SYSD.



Configuration Three

SYSC and SYSD do not share tape subsystems. SYSC uses tape addresses of 0380-039F. SYSD addresses its tape units with 0580-059F. Because these two systems use different addresses for different tape subsystems, the volume mount analyzer correctly interprets that the addresses are for different tape subsystems.

Configuration Four

SYSC and SYSD share a tape subsystem. However, each system addresses the same tape subsystem using a different address range. In this case, the volume mount analyzer incorrectly interprets the data as indicating SYSC and SYSD each use a different tape subsystem.

Merging the SMF data from SYSC and SYSD in Configuration Four might cause misleading or incorrect results for the Usage Report because the volume mount analyzer does not know that unit 0380 on SYSC and unit 0580 on SYSD are the same device. See "The Usage Report" on page 95 for more information. The volume mount analyzer assumes that each unique address represents a separate device, as shown in Configuration Three.

Note: Do not merge SMF data from multiple systems if your installation has Configuration Four. Analyze these SMF samples separately.

Avoiding Confusion with Shared Tape Subsystems

Refer to your system configuration maps or contact your systems programmer to find out what your system configuration is. If you have either Configuration One or Configuration Four, collect SMF records for each system and then run the volume mount analyzer separately for each system. Do not combine the SMF data from these systems before you run volume mount analyzer.

Knowing Notational Conventions

A uniform notation is used to describe the syntax of commands or the format of control records. This notation is not part of the language. The following conventions might be used in this book:

- [] Brackets enclose an optional entry. You may, but need not, include the entry. Examples are:
 - *[length]*
 - **[MF=E]**
- | An OR sign (a vertical bar) separates alternative entries. You must specify one, and only one, of the entries unless you allow an indicated default. Examples are:
 - **[REREAD | LEAVE]**
 - *[length | 'S']*
- { } Braces enclose alternative entries. You must use one, and only one, of the entries. Examples are:
 - **BFTEK={S | A}**
 - **{K | D}**
 - *{address | S | O}*

Sometimes alternative entries are shown in a vertical stack of braces. An example is:

```
MACRF={(R[C | P])}
{(W[C | P | L])}
{(R[C],W[C])}
```

In the example above, you must choose only one entry from the vertical stack.

- . . . An ellipsis indicates that the entry immediately preceding the ellipsis may be repeated. For example:
 - *(dcbaddr,[options],. . .)*
- ' ' A ' ' indicates that a blank (an empty space) must be present before the next parameter.

UPPERCASE BOLDFACE

Keywords are shown in **UPPERCASE BOLDFACE** in the syntax. These entries consist of keywords and punctuation symbols, such as commas, parentheses, and equal signs. Examples are:

- **CLOSE , , , ,TYPE=T**
- **MACRF=(PL,PTC)**

UNDERScoreD UPPERcase BOLDFACE

Indicates the default used if you do not specify any of the alternatives. Examples are:

- **[EROPT={ACC | SKP | ABE}]**
- **[BFALN={F | D}]**

Lowercase Italic

Lowercase italic indicates a value that you will supply according to specifications and limits described for each parameter. Examples are:

- *number*
- *image-id*
- *count*

Chapter 2. Using GFTAXTR to Extract Data from SMF Records

After you have collected SMF records for the selected period of time, you are ready to run GFTAXTR. This chapter describes the GFTAXTR program and shows you how to use it to extract data from SMF records. See “Collecting SMF Records” on page 4 for details.

It is important to know that because sort is invoked 20 or more times during processing, GFTAXTR requires DFSORT or its functional equivalent.

Understanding the GFTAXTR Program

The GFTAXTR program processes raw SMF data and transforms it into a format that GFTAVMA can analyze. The amount of SMF data is so large that it must stay on tape. Also, GFTAVMA reads the data many times, which would cause thousands of tape mounts. To make the repeated runs of GFTAVMA more efficient, GFTAXTR sorts and reduces the amount of data.

GFTAXTR Objectives

GFTAXTR has two objectives:

1. To reduce the quantity of input data so that GFTAVMA can be run iteratively using a DASD data set as input. Because the GFTAVMA analysis is an iterative process, using SMF records directly as input would be inefficient and costly in tape mounts.
2. To preprocess the SMF data by filling in missing record information and re-constructing the flow of the job that created each data set. GFTAXTR establishes all the inter-relationships between the data set and volume records, which enhances the performance of each GFTAVMA run.

GFTAXTR Processing

GFTAXTR extracts only tape-related records, which are the records GFTAVMA analyzes, reducing the size of the SMF input by approximately 95%. It also reduces each SMF record so as to contain only those fields necessary to perform tape activity analysis. GFTAXTR then provides missing information in the SMF records, such as invalid record format, data set name syntax errors, BLKSIZE value, and volume serial numbers.

Unlike GFTAVMA, which can be run many times using various keywords, you run GFTAXTR just once to create the input data set for GFTAVMA. The only other time you would run GFTAXTR is when you want to analyze a new sample of SMF data from another time period or another system.

The extract process was split, and several new modules were created. This allows you to make daily or weekly runs and then concatenate these runs into a final extract run. The output from this final run will be the input data set for GFTAVMA. However, if you wish, you can still use GFTAXTR to make just one run.

Note: Although the **TIME** and **DATE** keywords (see “Using GFTAXTR Keywords” on page 14) create subsets of the GFTAXTR data set, you cannot concatenate or combine multiple GFTAXTR output data sets into a single run without first enabling the concatenation code.

Understanding the Input to GFTAXTR

The GFTAXTR input is your collected SMF records. Since the order of SMF records on tape does not matter to the volume mount analyzer, these records can be in any chronological order.

When collecting SMF records, ensure that the SMFPRMxx member of SYS1.PARMLIB identifies the required record types for each system you want to analyze. See *z/OS MVS System Management Facilities (SMF)* for more information.

Required SMF Records

GFTAXTR input *must* contain SMF record types 14, 15, and 21. The input *must* also contain the data from record type 30 (sub-types 4 and 5) or record types 04, 05, 34 and 35.

If record types 4, 5, 21, 30 (sub-types 4 or 5), 34, or 35 do not match their corresponding type 14/15 records, GFTAXTR discards them.

Note: Record types 14 and 15 are mandatory input to GFTAXTR because they are the tape data set records. If you do not supply type 14 and 15 SMF records, or if GFTAXTR processing excludes all type 14 and 15 records through filtering, GFTAVMA processing terminates with a return code of 8.

GFTAXTR re-creates information from missing step end (type 4), job end (type 5), and demount (type 21) SMF records. When type 21 records are missing, GFTAXTR generates a type 21 record for every data set open record, which creates far more demounts than actually occurred. In some cases, type 4 and type 5 SMF records for long running jobs such as DFSMSHsm™, IMS/ESA®, or other started task jobs are not included in the collected SMF records. This is because the job or step might have started before the sample collection began or ended after the sample collection ended. GFTAXTR therefore re-creates these records.

Obviously, the degree of accuracy is much higher if the information for all record types is supplied in the SMF records instead of being re-created.

Although GFTAVMA will not terminate with errors if type 21 SMF records are missing, you must include type 21 records in the analysis because the information in type 21 records is critical to GFTAVMA’s ability to reconstruct the flow of jobs and to accurately report on tape mounts.

Note: *You must have type 21 records to do a tape analysis.* You cannot accurately study tape mounts without analyzing the type 21 SMF records. This is because without these records, the values in all of the volume mount analyzer reports are inaccurate.

SMF Record Types

Table 1 on page 11 shows SMF record types. See *z/OS MVS System Management Facilities (SMF)* for more information.

Table 1. SMF Input Records

Record Type	Record Description
04	Step End
05	Job End
14	EOV or CLOSE when open for reading. Called "open for input" in reports.
15	EOV or CLOSE when open for writing. Called "open for output" in reports.
21 ¹	Volume Demount
30 ²	Address Space Record (Contains sub-types 04, 05, 34, 35, and others)
34	Step End (TSO)
35	Job End (TSO)

Notes:

1. Type 21 records exist only for tape data.
2. Record type 30 (sub-types 4 and 5) is a shell record that contains the same information that is in record types 04, 05, 34, and 35.
If a type 30 record has the same data as type 04, 05, 34 and 35 records in the input data set, then GFTAXTR uses the data from the type 30 record and ignores the other records.

Submitting the GFTAXTR JCL

To submit the GFTAXTR job, you:

1. Copy the JCL member, GFTAXTRP, from SYS1.SAMPLIB
2. Change *only* the indicated fields
3. Submit the GFTAXTR JCL for execution using the TSO SUBMIT command

For more information on using JCL, see *z/OS MVS JCL Reference*.

Since this job runs for a long time, you might want to run it overnight.

GFTAXTR reads one input data set (SMFIN data set) and one input control statement data set (defined by the XTRCNTL DD statement) and builds a single output data set (XTRCIN data set).

```

//SMFXTRCT JOB (ACCT#,'ACCOUNTING-INFORMATION'),'PROGRAMMER NAME',
//      MSGCLASS=A,MSGLEVEL=(1,1),USER=&SYSUID,NOTIFY=&SYSUID
//*
//GFTAXTR PROC HLQ=,          * GFTAXTR FILE HIGH LEVEL QUALIFIER
//      RUN=,                * GFTAXTR FILE LOW LEVEL QUALIFIER
//      REGXTR=6144K,        * DEFAULT REGION FOR 'GFTAXTR'
//      SMFIN=,              * INPUT SMF DATA SET NAME
//      TAPES='20000,3000', * DEFAULT NUMBER OF TAPE MOUNTS 1
//      UNIT=SYSDA          * DEFAULT DASD ESOTERIC UNIT NAME
//*****
//*
//* GFTAXTR -- VOLUME MOUNT ANALYZER - SMF EXTRACTION UTILITY *
//*
//*****
//*
//* THIS IS THE "GFTAXTR" JCL.  USE "GFTAXTR" TO SELECT SMF RECORDS *
//* FOR INPUT TO "GFTAVMA".  "GFTAVMA" CANNOT TAKE SMF RECORDS *
//* DIRECTLY.  "GFTAVMA" INPUT MUST COME FROM "GFTAXTR". *
//*
//* "GFTAVMA" USES DATA FROM SMF RECORD TYPES 14, 15, 21, AND 30 *
//* (SUB-TYPES 4 AND 5).  IT WILL ALSO USE TYPES 4, 5, 34, AND 35 *
//* RECORDS AS AN ALTERNATIVE OR ANY COMBINATION OF ALL OF THE ABOVE.*
//*
//* PLEASE PERFORM THE FOLLOWING STEPS: *
//*
//* 1) MODIFY THE "JOB" STATEMENT TO FIT YOUR INSTALLATION'S *
//*    REQUIREMENTS. *
//*
//* 2) MODIFY THE "SMFIN" PARAMETER TO POINT TO THE EITHER THE *
//*    ORIGINAL SMF RECORDS OR TO AN SET OF RECORDS EXTRACTED *
//*    USING THE 'IFASMFDP' UTILITY. *
//*
//* 3) ALL OTHER 'GFTAXTR' PROC PARAMETERS WILL DEFAULT.  VERIFY *
//*    THAT THESE DEFAULTS ARE SATISFACTORY. *
//*
//* NOTE: *
//* TO SET THE 'TAPES' VALUE PROPERLY FOR CORRECT SPACE 1
//* ESTIMATES ON ALL DD STATEMENTS IN THIS PROC, USE THE *
//* NUMBER OF TYPE 21 RECORDS IN THE INPUT SMF SAMPLE AS *
//* THE VALUE.  THE 2ND VALUE SHOULD BE SPECIFIED AS A *
//* SAFETY VALVE FOR OVERFLOW.  FOR EXAMPLE, IF THE INPUT *
//* SAMPLE HAD 28,653 TYPE 21 RECORDS, ONE MIGHT SET THE *
//* TAPES='30000,5000'. *
//*
//*****

```

1 Use the TAPES parameter to do a symbolic override for the *allocation* of data sets. Refer to &TAPES in the SPACE parameter of the DD statement. Use the number of type 21 records from the IFASMFDP run as an estimate for the TAPES value on the PROC keyword. You can estimate the size of the SMF data extract file by looking at the printout from IFASMFDP. For example, if there were 21,732 type 21 records in the IFASMFDP output, you should set TAPES='22000,5000'.

See *z/OS MVS System Management Facilities (SMF)* for information on using the SMF dump utility, IFASMFDP.

```

/*
//DELETE EXEC PGM=IEFBRI4
//OLDXTRCT DD DISP=(MOD,DELETE),
//          DSN=&HLQ..GFTAXTR.&RUN,
//          UNIT=&UNIT,
//          SPACE=(TRK,(0))
/*
//XTRACT EXEC PGM=GFTAXTR,REGION=&REGXTR
//SMFIN DD DISP=SHR,DSN=&SMFIN
//SYSOUT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//XTRCIN DD UNIT=&UNIT,DSN=&HLQ..GFTAXTR.&RUN,          2
//          DISP=(NEW,CATLG),
//          RECFM=VB,
//          SPACE=(1000,(&TAPES),RLSE),AVGREC=U
//XTRCOUT DD UNIT=&UNIT,DSN=&XTRCOU,
//          RECFM=VB,
//          SPACE=(1000,(&TAPES),RLSE),AVGREC=U
//XTRCWK01 DD UNIT=&UNIT,
//          SPACE=(300,(&TAPES),,CONTIG,ROUND),AVGREC=U
//XTRCWK02 DD UNIT=&UNIT,
//          SPACE=(300,(&TAPES),,CONTIG,ROUND),AVGREC=U
//XTRCWK03 DD UNIT=&UNIT,
//          SPACE=(300,(&TAPES),,CONTIG,ROUND),AVGREC=U
//          PEND
//*****
/*
/* YOU MUST SET THE "SMFIN" PARAMETER VALUE TO THE NAME OF YOUR
/* SMF INPUT DATA SET. ALSO, SET "?HLQ" AND "?RUN" PARAMETERS
/* TO SPECIFY THE OUTPUT FROM THE 'GFTAXTR' RUN.
/*
/*
//*****
//          EXEC GFTAXTR,HLQ=?HLQ,RUN=?RUN,          3
//          SMFIN=?SMF.INPUT.DATASET.NAME'          4
/*
/*
//*****
/* THIS VERSION OF 'GFTAXTR' WILL RUN WITHOUT ANY INPUT KEYWORDS.
/* IF YOU DO SPECIFY 'GFTAXTR' KEYWORDS, THEN PLACE DIRECTLY
/* FOLLOWING THE "XTRCNTL" DD STATEMENT BELOW.
/*
//*****
//*****
//XTRACT.XTRCNTL DD *

```

6

2 The XTRCIN DD statement points to the GFTAXTR output data set that contains the SMF records that are input to GFTAVMA. Change the HLQ, RUN, and TAPES parameters with symbolic overrides on either the PROC statement in **1** or on the EXEC statement in **3**.

3 Set the substitutional parameters for the GFTAXTR PROC to your HLQ and RUN values. HLQ is the high-level qualifier of your SMF input data set. RUN is the low-level qualifier of your SMF input data set. Therefore, if you specified HLQ=ABC and RUN=RUN6, the name of the GFTAXTR output data set would be ABC.GFTAXTR.RUN6.

If you do not specify values for HLQ and RUN, you will get a JCL error.

4 Replace SMF.INPUT.DATASET.NAME with the name of your SMF data set. Your SMF input data set contains the set of SMF records collected by the IFASMFDP dump utility. This data set can contain many different types of SMF records, but GFTAXTR processes only the records necessary for tape analysis and discards all other records.

If you do not specify a value for SMFIN, you will get a JCL error.

If your SMF input sample is contained in GDGs, you can read the entire set of GDGs by specifying the GDG base name.

5 To specify your data set name on the EXEC statement, add the following after the EXEC statement:

```
//          EXEC GFTAXTR,SMFIN='your.dataset.name',HLQ=...,RUN=...
```

To include more than one input data set containing SMF data, add these statements after the EXEC statement:

```
//          EXEC GFTAXTR,SMFIN=,HLQ=...,RUN=...
//XTRACT.SMFIN DD DSN=your.data.set1,DISP=SHR
//          DD DSN=your.data.set2,DISP=SHR
//          DD DSN=your.data.set3,DISP=SHR
```

These statements nullify the SMFIN statement on the EXEC statement, and a JCL override DD statement, XTRACT.SMFIN, refers to additional data sets.

6 Specify optional keywords following the XTRACT.XTRCNTL DD statement.

There are no required keywords for GFTAXTR. However, you should use the DATE keyword to select the one month of your SMF sample you want to study.

See “Using GFTAXTR Keywords” for the rules and syntax of the GFTAXTR keywords.

Submitting the GFTAXTR JCL for Concatenated Input

To submit the GFTAXTR job to concatenate input, you:

1. Copy the JCL member, GFTAXTRP, from SYS1.SAMPLIB
2. Change all references from GFTAXTR to GFTAXTD when making daily or weekly runs, and to GFTAXTF when making the final run. Then, change only the indicated fields
3. Submit the GFTAXTD or GFTAXTF JCL for execution using the TSO SUBMIT command

For more information on using JCL, see *z/OS MVS JCL Reference*.

GFTAXTD reads one input data set (SMFIN data set) and one input control statement data set (defined by the XTRCNTL DD statement) and builds a single output data set (XTRCIN data set) which can then be concatenated with other GFTAXTD runs to become the input for the final run (GFTAXTF). Make sure that the output from each GFTAXTD run has a unique name, such as RUN=RUN1, RUN2.

GFTAXTF reads one input data set (SMFIN data set) and one input control statement data set (defined by the XTRCNTL DD statement) and builds a single output data set (XTRCIN data set) which is used as input to GFTAVMA. To include all GFTAXTD data sets for the final run, add these statements after the EXEC statement:

```
//          EXEC GFTAXTF,SMFIN=,HLQ=...,RUN=...
//XTRACT.SMFIN DD DSN=your.data.set1,DISP=SHR
//          DD DSN=your.data.set2,DISP=SHR
//          DD DSN=your.data.set3,DISP=SHR
```

These statements nullify the SMFIN statement on the EXEC statement, and a JCL override DD statement, XTRACT.SMFIN, refers to additional data sets.

Using GFTAXTR Keywords

There are four optional keywords that you can use in the GFTAXTR JCL. You can use GFTAXTR keywords to select SMF records by date and time, identify the field for accounting information, and control the printing of the output report.

Keyword Summary and Syntax

The XTRCNTL DD statement points to a sequential data set with 80-byte records that contain GFTAXTR keywords. Normally, this is a DD * data set.

As pointed out earlier, all GFTAXTR keywords are optional. If you do not want to use any keywords, comment out the XTRACT.XTRCNTL DD statement. Table 2 summarizes the keywords for GFTAXTR.

Table 2. Keyword Summary for GFTAXTR

[ACCTFLD (<i>nn</i>)]
[DATE (<i>fromdate</i> [, <i>todate</i>])]
[LINES (<i>nn</i>)]
[TIME (<i>fromtime</i> [, <i>totime</i>])]

Keyword Rules

When using GFTAXTR keywords,

- Start or end the keywords in any position from 1 to 71. Positions 72 through 80 are ignored.
- Do not use continuation characters. GFTAXTR does not allow continuation characters.

Comments are surrounded by the delimiters /* and */. A comment start delimiter, /*, cannot be in column one.

ACCTFLD Keyword

Use the ACCTFLD keyword as an index to the accounting information that is specified in the JOB statement of those jobs that created or used tape. The ACCTFLD value indexes to the SMF5ACTF field of the type 5 SMF record and selects the next eight characters of accounting information. GFTAVMA uses these eight characters as the job's accounting information.

Abbreviation: AFLD

Syntax

Figure 1 shows the syntax for the ACCTFLD keyword.

[ACCTFLD (<i>nn</i>)]

Figure 1. ACCTFLD Keyword Syntax

nn Specifies the number that GFTAXTR uses to index-by-eight into the accounting information and select eight characters to use as the accounting data in the GFTAVMA program.

You can specify a value from 1 to 18.

Default

If you do not specify the ACCTFLD keyword, GFTAXTR uses the default of 1, which is the first eight characters of accounting information.

Relationship to Other Keywords

ACCTFLD indexes into the accounting information filtered by the GFTAVMA ACCOUNT filtering keyword. See “ACCOUNT Keyword” on page 64 for more information.

Reports Affected

ACCTFLD affects the TOP(ACCOUNT) report. See “The Top Report” on page 92 and “Parameters for the Top Report” on page 149 for more information.

DATE Keyword

Use **DATE** to select SMF records created on specified dates or between specified dates if you request a range.

Abbreviation: None

Syntax

Figure 2 shows the syntax for the **DATE** keyword.

[DATE (*fromdate*[,*todate*])]

Figure 2. DATE Keyword Syntax

fromdate

Specifies the start date of the SMF collection period.

todate

Specifies the end date of the SMF collection period.

If you specify only a *fromdate*, GFTAVMA selects only the records created on that date.

Use the **DATE** keyword to select the dates for the period you want to study. The time defaults to 00:00:00 to 23:59:59.

You can use any of the following formats for *fromdate* and *todate*:

mm/dd/yy	The month/day/year specified with 2 digit fields, for example: 03/15/93
mm/dd/yyyy	The month/day/year where the month and day are 2 digit fields and the year is a 4 digit field, for example: 03/15/2004
yyyddd	The year and Julian date using 4 digits for the year, for example: 2004074
yyddd	The year and Julian date using 2 digits for the year, for example: 93074

Default

If you do not specify the **DATE** keyword, GFTAXTR analyzes the SMF records for every day in the collected SMF sample.

DATE Examples

If you specify:

DATE (05/30/93)

GFTAXTR only includes the SMF records created on May 30, 2004.

If you specify:

DATE (05/27/93,06/28/93)

GFTAXTR includes the SMF records created from May 27, 2004, through June 28, 2004.

LINES Keyword

Use **LINES** to set the number of lines per page for the printed output report.

Abbreviation: None

Syntax

Figure 3 shows the syntax for the **LINES** keyword.

[LINES (*nn*)]

Figure 3. LINES Keyword Syntax

nn Represents the number of output lines to print per page. The value must be a whole number between 9 and 99.

Default

The default for **LINES** is 60 lines per page.

TIME Keyword

Use **TIME** to select records based on the timestamp when the SMF record was created. You can also use **TIME** to examine a daily processing window or a particular shift.

If you specify a time range using **TIME** (*fromtime,totime*), only records in that time range are extracted for each day in the sample.

Abbreviation: None

Syntax

Figure 4 shows the syntax for the **TIME** keyword.

[TIME (*fromtime*[,*totime*])]

Figure 4. TIME Keyword Syntax

Both *fromtime* and *totime* have this syntax:

hh:mm:ss

Specifies the hour, minute, and second of the timestamp.

If you specify only a *fromtime*, the default *totime* is the end of the day, 23:59:59.

The valid range of values is 00:00:00 for midnight, which is the start of the day, and 23:59:59, which is the end of the day.

Times can be specified as hh:mm:ss, hh:mm, or just hh. The specified time values must be in 24 hour clock values.

TIME

Default

If you do not specify the **TIME** keyword, GFTAXTR analyzes the SMF records for all times in the collected SMF sample.

TIME Examples

If you specify:

TIME (3)

GFTAXTR includes only records created between 03:00:00 and 23:59:59 on each day of the sample period.

If you specify:

TIME (8,12)

GFTAXTR includes only records created between 08:00:00 and 12:00:00 on each day of the sample period.

Interpreting the GFTAXTR Output Report

GFTAXTR produces a single output report in the SYSPRINT data set. This report shows the results of the various phases of SMF record processing and contains all GFTAXTR error and informational messages. See Appendix D, “Volume Mount Analyzer Messages,” on page 167 for more information.

The GFTAXTR output report is informational only. It is *not* the input to GFTAVMA. The input to the GFTAVMA program is the output data set defined by the XTRCIN DD statement in the GFTAXTRP JCL.

The GFTAXTR output report shows the SMF records that were dropped, created, read, and written during the processing phases.

The output report from GFTAXTD daily or weekly runs will only contain the internal record generation phase. The output report from the GFTAXTF final run will contain all phases of the SMF record processing.

GFTAXTR Output Report—Input Phase

If you specify keywords in the input phase, GFTAXTR filters the SMF records.

Figure 5 is an example of part of the GFTAXTR output report that shows what keywords were used to filter the SMF records for this run.

```
THE FOLLOWING KEYWORDS WERE INPUT TO THIS RUN:  
DATE(05/27/2002,06/28/2002) 1
```

```
ALL INPUT SMF RECORDS WILL BE FILTERED BY THE FOLLOWING DATE/TIME STAMP RANGE:  
DATE RANGE: 05/27/2002 TO 06/28/2002  
TIME RANGE: 12:00:00 A.M. TO 11:59:59 P.M.
```

Figure 5. Example GFTAXTR Report—Input Phase

1 The **DATE** keyword was specified for this run. By using the **DATE** keyword, you can specify the end date of your sample. Since **TIME** was not specified, GFTAXTR analyzed all SMF records for all times during the specified date range.

GFTAXTR Output Report—Messages

During the internal record generation phase, GFTAXTR reads, drops, corrects, and writes SMF records. GFTAXTR also generates messages that describe why records were dropped and provide other processing information.

Figure 6 is an example of part of the GFTAXTR output report that shows the volume mount analyzer messages and the records GFTAXTR dropped and corrected during this processing phase.

```
GFTAXTR -- INTERNAL RECORD GENERATION PHASE -- 05/21/2003 -- 02:33:02 P.M.

GFTA077W SMF INPUT ERROR -- THE FOLLOWING DEMOUNT (TYPE 21) RECORD WAS MISSING THE VOLUME
INFORMATION -- RECORD WAS DROPPED
-- JOB= (          - 12:00:00 A.M.) - UADR=0520 - TSMV=05/28/2002 - 08:11:14 A.M.

Many other messages here 2
    other messages
    other messages ...

SMF DROPPED RECORD ERROR SUMMARY:
  132  DATA SET NAME ERRORS ('SMFJFCB1' FIELD)
  140  VOLUME SERIAL ERRORS ('SMF21VOL' FIELD)
-----
  272  SMF RECORD   ERRORS (TOTAL)

SMF CORRECTED FIELD ERROR SUMMARY:
   13  DSNAME  FIELD ERRORS (INVALID CHARACTERS)
-----
   13  SMF    FIELD ERRORS CORRECTED (TOTAL)
```

Figure 6. Example GFTAXTR Report—Internal Record Generation Phase (Messages)

2 This report usually includes a series of messages that describe errors GFTAXTR detected during the processing of SMF records.

GFTAXTR drops an SMF record because:

- The record has an error, such as a character field contains binary data or a binary field contains character data.
- The record is incomplete.

SMF records might contain conflicting or incomplete pieces of data. GFTAXTR often compensates for incomplete data by reconstructing the flow of the job from other records.

Many of the messages relating to created or corrected records are informational only, but they show the extent to which the volume mount analyzer has used its own logic to provide missing information in the SMF records. See Appendix D, “Volume Mount Analyzer Messages,” on page 167 for the complete messages.

GFTAXTR Output Report—Internal Record Generation Phase

In this phase, GFTAXTR takes the raw SMF records, discards the parts of the records GFTAVMA does not use, and builds new records in a compressed format.

Figure 7 is an example of part of the GFTAXTR output report that shows records read, generated, dropped, generated and written while GFTAXTR was converting the SMF data into input for the GFTAVMA program.

```

GFTAXTR -- INTERNAL RECORD GENERATION PHASE -- 05/21/2003 -- 02:33:02 P.M.
RECORD ANALYSIS PERIOD:
                                START: 05/27/2002 -- 12:02:01 A.M. 3
                                END:   06/28/2002 -- 11:59:18 P.M.

      0 TYPE 04 RECORDS READ (STEP TERMINATION)
      0 TYPE 05 RECORDS READ (JOB TERMINATION)
    20993 TYPE 14 RECORDS READ (OPEN FOR INPUT)
    82899 TYPE 15 RECORDS READ (OPEN FOR OUTPUT)
    48669 TYPE 21 RECORDS READ (TAPE DEMOUNTS)
    76622 TYPE 30 RECORDS READ (COMMON ADDRESS SPACE RECORDS)
      0 TYPE 34 RECORDS READ (TIME SHARING TERMINATION)
      0 TYPE 35 RECORDS READ (TSO LOGOFF)
-----
    229183 TYPE XX RECORDS READ (TOTAL)
      0 TYPE XX RECORDS IGNORED (OTHER RECORD TYPES)
-----
    229183 TYPE XX RECORDS PROCESSED (TOTAL FROM 'SMFIN' INPUT FILE) 4

      0 TYPE 14 RECORDS DROPPED (BY DEVICE TYPE MISMATCH) 5
    3482 TYPE 14 RECORDS DROPPED (DATE/TIME FILTERING) 6
     132 TYPE 15 RECORDS DROPPED (SMF RECORD ERRORS) 7
    15524 TYPE 15 RECORDS DROPPED (DATE/TIME FILTERING)
     140 TYPE 21 RECORDS DROPPED (SMF RECORD ERRORS)
     651 TYPE 21 RECORDS DROPPED (NO ASSOCIATED TYPE 14 OR 15 RECORD FOUND) 8
     9182 TYPE 21 RECORDS DROPPED (DATE/TIME FILTERING)
      0 TYPE 21 RECORDS CREATED (FOR UNMATCHED 14/15 RECORDS) 9
    14098 TYPE 30 RECORDS DROPPED (DATE/TIME FILTERING)
      0 TYPE 30 RECORDS DROPPED (UNUSED) 10
-----
    43209 TYPE XX RECORDS DROPPED (TOTAL)

    49546 TYPE 04 RECORDS GENERATED (FROM TYPE 30 RECORDS) 11
    12978 TYPE 05 RECORDS GENERATED (FROM TYPE 30 RECORDS)
    17511 TYPE 14 RECORDS WRITTEN (FROM TYPE 14 RECORDS)
    67243 TYPE 15 RECORDS WRITTEN (FROM TYPE 15 RECORDS)
    38696 TYPE 21 RECORDS WRITTEN (FROM TYPE 21 RECORDS)
-----
    185974 TYPE XX RECORDS WRITTEN ('XTRCIN' FILE) 12

```

Figure 7. Example GFTAXTR Report—Internal Record Generation Phase

- 3** Time/Date range of the input SMF sample.
- 4** Raw SMF records from the SMFIN file, which is the file of SMF records input to GFTAXTR.
- 5** Records dropped because they are not tape-related. Instead, they are DASD 14/15 records.
- 6** Records dropped because of date/time filtering.
- 7** Records dropped because of incomplete SMF records. Some SMF records are missing pieces of information that the volume mount analyzer needs, such as a volume serial number or unit address.
- 8** Type 21 records dropped because they did not have matching type 14 or 15 SMF records. The type 21 records might not have matched type 14 or 15 records because the SMF interval sample did not go back far enough in time to include the matching 14/15 records. To GFTAXTR, the record mismatch was an error, but it was not an SMF error. Since GFTAXTR could find the related data set activity, it discarded the demount record.
- 9** Demount records (type 21) created because they did not have corresponding data set open records.

10 Type 30 SMF records dropped because they did not have a sub-type 04 or 05 record.

11 Records generated. GFTAXTR created these records with information pulled from other records types.

12 The total number of SMF records written to the output data set, XTRCIN. This number is smaller than that of SMF records read in **4**. The number of records read (229,183) minus the number of records dropped because of filtering and errors (43,209) equals the number of records in the output file (185,974).

The percentage of records dropped because of errors is small: 42,286 records were dropped by DATE/TIME filtering, 651 records were dropped by SMF sample interval, and 272 records were dropped because of SMF errors.

By the end of the internal record generation phase, GFTAXTR has cleaned up and reduced the SMF records.

GFTAXTR Output Report—Internal SMF Job/Step Reconstruction Phase

During this phase, GFTAXTR completely reconstructs the jobs that used tapes, including job steps, opens, and demounts.

Figure 8 on page 22 is an example of part of the GFTAXTR output report that shows the reconstruction of jobs with incomplete SMF data and the tape usage during the SMF data collection period.

TIME

GFTAXTR -- INTERNAL SMF JOB/STEP RECONSTRUCTION PHASE -- 05/21/2003 -- 02:35:41 P.M.

GFTA104W SMF INPUT ERROR -- DEMOUNT RECORD HAS NO MATCHING 14/15 RECORD -- THE FOLLOWING DEMOUNT
(TYPE 21) RECORD WAS DROPPED:

--JOB=BASN02D (06/06/2002 - 05:32:28 P.M.) - VOL=703772 - UADR=052F - TSMV=06/06/2002 - 06:39:03 P.M. **13**

GFTAXTR -- INTERNAL SMF JOB/STEP RECONSTRUCTION PHASE -- 05/21/2003 -- 02:35:41 P.M.

JOB STATISTICS:

116	TYPE 21 RECORDS WERE DROPPED BY 'GFTAXTR'	
6	TYPE 21 RECORDS WERE CREATED BY 'GFTAXTR'	14
5	INTERNAL STEP END(S) CREATED BY 'GFTAXTR'	15
2435	STEP END RECORD(S) MODIFIED BY 'GFTAXTR'	
5	INTERNAL JOB END(S) CREATED BY 'GFTAXTR'	
202	JOB END RECORD(S) MODIFIED BY 'GFTAXTR'	16
359	NON-TAPE JOBS PROCESSED	
12448	TAPE JOBS PROCESSED	17

12807	TOTAL JOBS PROCESSED	

OUTPUT RECORD STATISTICS ('XTRCIN' FILE):

17511	OPEN (INPUT) FOR TAPE	
67243	OPEN (OUTPUT) FOR TAPE	
38586	DEMOUNT RECORDS (TAPE)	

123340	TOTAL OUTPUT RECORDS WRITTEN	18

TAPE MOUNT STATISTICS:

10679	INPUT MOUNTS	
26280	OUTPUT MOUNTS	19
1627	INPUT/OUTPUT MOUNTS	

38586	DEMOUNT RECORDS (TAPE)	

Figure 8. Example GFTAXTR Report—Internal SMF JOB/STEP Reconstruction Phase

Note: The output report generated in Figure 8 is not the same as the output report generated in Figure 7 on page 20

13 This message tells you that a type 21 SMF record had no matching 14/15 SMF record and, therefore, GFTAXTR dropped the type 21 SMF record. See Appendix D, “Volume Mount Analyzer Messages,” on page 167 for more information.

14 Type 21 records that GFTAXTR created to replace missing records. GFTAXTR does not create type 14 or 15 records.

15 Internal step end records that GFTAXTR created to replace missing SMF records. The step end probably occurred outside of the date or time period of the sample.

16 Job end records that GFTAXTR modified to correct some timestamps. For example, the timestamp of a demount might be later than the timestamp of a corresponding job end. In cases like this, GFTAXTR modifies the timestamp of the job end to match the timestamp of the tape demount.

17 The number of tape and non-tape jobs in the SMF sample. GFTAVMA only analyzes tape data sets and the corresponding jobs that processed them.

18 The total number of output records written. These records are in an internal format specific to GFTAXTR.

19 Output mounts are those resulted in open for output activity *only*. This number will match the number of scratch mounts if all output activity required scratch tapes or tapes with no specific volume serials. A scratch mount is always an output mount, but the reverse is not necessarily true.

Normally, only the input mounts are specific mounts, and all other mounts would be scratch mounts. If you use specific volume serials for new tape allocations, those mounts would be specific as well. Most installations use a generic tape unit with no volume serial.

GFTAXTR Output Report—Data Set Statistics Phase

In this phase, GFTAXTR identifies multivolume and multifile data sets, classifies data set usage types (such as BCOPY, *backup*, and *active*), determines attributes (such as DS type, latest version, and total size), and prints these statistics. The term “multifile data sets” is defined in “Single Volume and Multifile” on page 29.

See “Data Set Usage Types” on page 36 for a description of how the volume mount analyzer determines data set usage types.

Figure 9 is an example of part of the GFTAXTR output report that shows the breakdown of non-DFSMSHsm tape data set types.

```
GFTAXTR -- DATA SET STATISTICS PHASE -- 05/21/2003 -- 02:42:35 P.M.
```

```
NON-HSM TAPE DATA SET TYPE BREAKDOWN: 20
```

```
38702 GDS DATA SET(S) -- 2049 GDG GROUP(S)
  18 TSO DATA SET(S)
  45 SYSTEM DATA SET(S)
 142 TEMP DATA SET(S)
2331 OTHER DATA SET(S)
-----
41238 TOTAL DATA SET(S)
```

```
NON-HSM TAPE DATA SET ALLOCATION STATISTICS:
```

```
38702 DATA SET(S) ( 93.9%) THAT ARE GENERATIONS
38956 DATA SET(S) ( 94.5%) THAT ARE CATALOGED 21
  125 DATA SET(S) ( 0.3%) ALLOCATED WITH SPACE 22
39469 DATA SET(S) ( 95.7%) WITH EXPIRATION DATE CODED 23
3690 DATA SET(S) ( 8.9%) REFERENCED WITHIN A SINGLE DAY 24
  90 DATA SET(S) ( 0.2%) WITH ZERO BLKSIZE IN SMF RECORDS 25
39365 DATA SET(S) ( 95.5%) WITH BLOCKED RECORDS
31834 DATA SET(S) ( 77.2%) WITH BLKSIZE > 1/2 TRK (3380) 26
20645 DATA SET(S) ( 50.1%) WITH BLKSIZE > 1/2 TRK (3390)
```

Figure 9. Example GFTAXTR Report—Data Set Statistics Phase

20 GFTAXTR determines data set types based on standard data set name formats.

21 The JCL for the jobs creating these data sets included `DISP=(,CATLG)`.

22 The JCL for the jobs creating these data sets included `SPACE=(x,(n))`.

Note: Although the `SPACE` parameter might have been coded on the DD statements of these data sets, their size might not be accurate because tape allocation does not use the `SPACE` parameter. If you use the specified `SPACE` values, there is no guarantee that these values are adequate for a DASD allocation.

23 Data sets with expiration dates. `EXPDT` is frequently used by tape management systems to code various tape management attributes.

24 Data sets (other than HSM, temporary, or backup) where every reference occurred on a single calendar day.

25 Data sets with a block size of zero. Unless you specify a value on the `BLKSIZE` keyword, GFTAVMA assumes a default of 32 760 for these data sets. See “`BLKSIZE` Keyword” on page 132 for details.

26 Data sets currently blocked at greater than a half track of 3380 or at greater than a half track of 3390, which causes inefficient use of DASD space.

You should address block size efficiency when planning for *migration*. An efficient `BLKSIZE` specification for tape might not be a good value to use when placing the data set on DASD. The amount of DASD space required is affected by the user-specified or system-determined `BLKSIZE` value.

GFTAXTR Output Report—NON-HSM Tape Data Set Usage Analysis

Figure 10 on page 25 is an example of the part of the GFTAXTR output report that shows the detailed data set usage for non-HSM tape data sets. See “Data Set Usage Types” on page 36 for the order of precedence that the volume mount analyzer uses to determine usage types.

GFTAXTR -- DATA SET STATISTICS PHASE -- 05/21/2003 -- 02:42:35 P.M.

NON-HSM TAPE DATA SET USAGE ANALYSIS:

```

2668  MULTIVOLUME DATA SETS ( 6.5%)
-     126 'USAGE(BACKUP)' DATA SETS ( 0.3%)
-     4 MULTIFILE DATA SETS
-     122 SINGLE-FILE DATA SETS
-     32 'USAGE(BCOPY)' DATA SETS ( 0.1%)
-     32 MULTIFILE DATA SETS
-     0 SINGLE-FILE DATA SETS
-     171 'USAGE(ACTIVE)' DATA SETS ( 0.4%)
-     65 MULTIFILE DATA SETS
-     106 SINGLE-FILE DATA SETS
-     2295 'USAGE(SINGLE)' DATA SETS ( 5.6%)
-     2096 MULTIFILE DATA SETS
-     199 SINGLE-FILE DATA SETS
-     44 'USAGE(TEMP)' DATA SETS ( 0.1%)
-     2 MULTIFILE DATA SETS
-     42 SINGLE-FILE DATA SETS
38570 SINGLE-VOLUME DATA SETS ( 93.5%)
-     894 'USAGE(BACKUP)' DATA SETS ( 2.2%)
-     683 MULTIFILE DATA SETS
-     211 SINGLE-FILE DATA SETS
-     674 'USAGE(BCOPY)' DATA SETS ( 1.6%)
-     379 MULTIFILE DATA SETS
-     295 SINGLE-FILE DATA SETS
-     1431 'USAGE(ACTIVE)' DATA SETS ( 3.5%)
-     581 MULTIFILE DATA SETS
-     850 SINGLE-FILE DATA SETS
-     35474 'USAGE(SINGLE)' DATA SETS ( 86.0%)
-     27439 MULTIFILE DATA SETS
-     8035 SINGLE-FILE DATA SETS
-     97 'USAGE(TEMP)' DATA SETS ( 0.2%)
-     14 MULTIFILE DATA SETS
-     83 SINGLE-FILE DATA SETS

```

34

Figure 10. Example GFTAXTR Report—Non-HSM Tape Data Set Usage Analysis

27 Multivolume data sets. A multivolume data set extends across volume boundaries; it is part of a group of data sets that collectively spans volume boundaries. See “Single/Multi Volume/File Data Sets” on page 29 for more information.

28 Backup data sets. Since these data sets were opened for output only, the volume mount analyzer categorized them as backup. They might not all be backup, but most usually are.

29 Active data sets. Since these data sets were either opened for input only or have a mix of input and output activity, the volume mount analyzer categorized them as active.

30 Single data sets. Since these data sets were referenced in the input sample within a single calendar day, the volume mount analyzer categorized them as single.

31 Temporary data sets. Temporary data sets exist only for the life of the job. You might want to eliminate these data sets from tape and store in a DASD buffer until they expire.

32 Single-volume data sets. A single-volume data set does not extend across volume boundaries. It often resides alone on a volume and is usually small. Having many small, single-volume data sets might reduce the efficiency of your tape media usage. See “Single/Multi Volume/File Data Sets” on page 29 for more information.

33 Backup copy of data sets. Since these data sets were opened once for input and once for output, the volume mount analyzer categorized them as BCOPI.

34 An output sample that does not show DFSMSHsm output statistics. When output includes DFSMSHsm statistics, the report shows DFSMSHsm single file data sets, which are the large multivolume data sets into which DFSMSHsm writes user data sets in blocks of 16384 bytes.

There are two single-file data sets for migration and two for backup. GFTAVMA generates two data sets for each to distinguish inputs from outputs in the analysis. GFTAVMA uses an appended LLQ qualifier of .INPUT and .OUTPUT for each of the single file format (SFF) data sets to accumulate statistics separately for the data that DFSMSHsm is writing versus the data DFSMSHsm is reading. In GFTAVMA reports, each of the SFF data sets is treated as two data sets. See “Data Set Usage Types” on page 36 for more information about DFSMSHsm processing.

GFTAXTR Output Report—Final Record Output Phase

Figure 11 is an example of part of the GFTAXTR output report that shows SMF record output statistics. These are the records that are input to GFTAVMA.

```
GFTAXTR -- FINAL RECORD OUTPUT PHASE -- 05/21/2003 -- 02:57:22 P.M.

JES SYSTEM ID(S):

      SYSTEM ID   OPER SYS      35
      -----   -
      SYSA       MVS/ESA

DATA COMPRESSION STATISTICS:      36

120619.6 KB (TOTAL IN THE SMF INPUT FILE)
30452.0 KB (IN THE GFTAXTR OUTPUT FILE)
25.2% (DATA COMPRESSION RATIO)

OUTPUT RECORD STATISTICS:

      1 TYPE 000 RECORD(S) (GFTAXTR VERSION RECORD)
      1 TYPE 001 RECORD(S) (JES SYSTEM ID RECORD)
17511 TYPE 014 RECORD(S) (DATA SET OPEN (INPUT)) 37
67243 TYPE 015 RECORD(S) (DATA SET OPEN (OUTPUT))
38586 TYPE 021 RECORD(S) (TAPE VOLUME DEMOUNT)
-----
123342 TYPE XXX RECORD(S) (TOTAL IN OUTPUT FILE) 38

GFTAXTR PROGRAM COMPLETED AT 02:57:48 P.M. ON 05/21/2003 -- RETURN CODE = 4 39
```

Figure 11. Example GFTAXTR Report—Final Record Output Phase

35 This shows the system identifiers for the systems you are studying in the SMF sample. Systems in a run must all belong to the same logical complex with shared access to the tape subsystems being analyzed. See “Understanding the Limitations of the Volume Mount Analyzer” on page 5 for shared systems considerations.

After you run GFTAVMA, make sure that the system IDs and the date and time on this GFTAXTR report match the system IDs and date and time on the GFTASRT1 summary report. See “GFTASRT1—Input Analysis Report” on page 52.

36 GFTAXTR has taken the original SMF data, extracted only the parts necessary for the analysis, and compressed the data. These statistics reflect the efficiency of this process. In this example, the output from GFTAXTR uses only 25.2% of the storage that the original SMF data would have used.

37 Data set open records do not always correlate with previous processing phases because in some cases a data set can have more than one open record. See “Open Counts” on page 34 for more information.

38 GFTAXTR creates its own internal record types. The counts of these records match the counts of the similar type SMF input records that were not dropped. But the format of these records is different from that of the SMF records. GFTAXTR has indexed data set and demount records to one another and added job flow information so that GFTAVMA can get an accurate picture of the jobs that created the data sets.

Because of GFTAXTR’s data compression, the total number of output records is smaller than the total in Figure 7 on page 20.

39 Ideally, the return code is 0. However, the return code is usually 4 because GFTAXTR generates warning messages to tell you the time and reasons that SMF records were dropped or altered. Missing information in the input SMF records caused these warning messages. Normally, the missing information only represents a small percentage of the total SMF input and does not affect the analysis. See Appendix D, “Volume Mount Analyzer Messages,” on page 167.

Examining the GFTAXTR Output

You should check your GFTAXTR output by asking the following questions:

- What percentages of the records were dropped and created?

Large percentages of 14/15 records are dropped because they are DASD records, so they are not studied in a tape analysis. The number of 14/15 records dropped might actually be larger than the number retained.

If records have a time mismatch, the volume mount analyzer creates type 21 records for the 14/15 records that do not have matching type 21 records. If things run normal, then only 1% to 2% of the records created will be type 21 records. However, if a large number of type 21 records are created, this means there were many time mismatches.
- Were type 21 SMF records included in the study?

You must have type 21 records for tape analysis.
- What percentage of the data sets was USAGE (BACKUP) or USAGE (SINGLE)?

The higher the percentage, the better.
- What percentage of data sets was allocated with SPACE?

This percentage should be low. Since tape allocation does not require a SPACE value, be aware of data sets that were allocated with SPACE. If you intercept these data sets, the SPACE value specified in the JCL might be incorrect.
- Did the return code print out at the end of the report?

You should either get a return code of 0 indicating no errors or 4 indicating some recoverable errors. Since there are usually SMF record time mismatches, the return code is often 4.

Troubleshooting when Obtaining the GFTAXTR Output

If you did not obtain the GFTAXTR output, make sure that:

- Your SMF input includes types 14 and 15 records.
- These 14 and 15 records are for tape, not DASD.
- You have type 21 records in your SMF input.
- You have type 30 records that correspond to tape mount and open activity.
- Tapes were mounted during the period specified for SMF record collection.
- You do not have many mismatches in the time frame of records for demounts (21) and records for opens (14/15).
- Your SMF records are not altered. They must correspond to the formats in *z/OS MVS System Management Facilities (SMF)*.
- Your SMF records have the correct length. If your installation writes its own SMF records, they might have incorrect length. GFTAXTR inspects each record type for a minimum length value. If the length of the record is less than the standard SMF record format, then GFTAXTR will not use the short records. See *z/OS MVS System Management Facilities (SMF)* for more information.
- You did not over-filter and exclude *all* SMF records.

Chapter 3. Understanding Tape—Some Concepts

This chapter describes the following tape concepts:


- single/multi volume/file data sets
- statistical mounts
- statistical volumes
- open counts
- tape cartridge usage
- data set usage types
- tape mount management data set categories

You need to understand these tape concepts in order to use GFTAVMA efficiently.

Single/Multi Volume/File Data Sets


To understand how the volume mount analyzer calculates mounts and tape volume usage, you must understand the four ways data sets can reside on tape volumes.

Single Volume and Single File




A single-volume, single-file data set resides on one volume and does not share that volume with other data sets.

Single Volume and Multifile



A single-volume, multifile data set resides on one volume, but shares that volume with at least one other data set.

Multivolume and Single File



A multivolume, single-file data set is one data set that spans two or more volumes. None of these volumes is shared with other data sets.

Multivolume and Multifile

VOLUME 1	A.B.C	D.E.F(1)
----------	-------	----------

VOLUME 2	D.E.F(2)	X.Y.Z
----------	----------	-------

A multivolume, multifile data set is a part of a collection of data sets that span two or more volumes. The name *multivolume*, *multifile* applies to the entire collection of data sets written together, whether each data set actually spans a volume boundary. In other z/OS manuals, this collection of volumes is called an aggregate.

In this example, A.B.C, D.E.F, and X.Y.Z are all multivolume, multifile data sets. A.B.C. is file 1; D.E.F is file 2, regardless of the volume; and X.Y.Z is file 3.

Statistical Mounts

By definition, mounts are associated with volumes. If a volume contains only one data set (single-file, single-volume), the number of mounts is directly associated with that data set in a ratio of one to one. Thus, for single-file data sets, the number of statistical mounts equals the number of mounts of the associated volumes.

In cases where data sets are stacked on a tape or across multiple tapes (multifile, multivolume), calculating mounts becomes more complicated because more than one data set can be written after a single mount.

Calculating statistical mounts for multifile, multivolume data sets means assigning fractions of mounts when *filtering* causes GFTAVMA to divide the mounts among several data sets on the volume. For multifile data sets, the number of statistical mounts is an average of the number of mounts for the volume divided equally among all data sets on that volume (statistical mounts = number of mounts/number of data sets on a volume).

Statistical mounts reflect all the mounts attributed to a data set during the reporting period. It is a cumulative value. On the other hand, *statistical volumes* reflect only the volumes that contain the latest version of the data set. See "Statistical Volumes" on page 31 for information on statistical volumes.

Example of Calculating Statistical Mounts

VOLUME 111111	A.B.C
---------------	-------

Volume 111111 was mounted 3 times.

Statistical mounts for A.B.C
= volume mounts (3)/data sets on volume (1)
= 3.0 statistical mounts

Data set A.B.C is a single-file data set; therefore, the number of statistical mounts (3.0) equals the number of mounts for volume 111111 (3).

Example for Multifile Data Sets

VOLUME 222222	A.B.C	X.Y.Z
---------------	-------	-------

Volume 222222 was mounted 3 times.

Statistical mounts for A.B.C
= volume mounts (3)/data sets on volume (2)
= 1.5 statistical mounts per data set

Data sets A.B.C and X.Y.Z share Volume 222222. For these multifile data sets, the number of statistical mounts is an average of the number of mounts for the volume divided equally among all data sets on that volume (statistical mounts = number of mounts/number of data sets on a volume). So a volume mounted three times containing two data sets (3/2) equals 1.5 statistical mounts.

Example for Multiple Versions of a Data Set

VOLUME 333333	A.B.C	X.Y.Z
---------------	-------	-------

VOLUME 444444	A.B.C'
---------------	--------

Volume 333333 was mounted 3 times.

Volume 444444, which contains a later version of A.B.C, was mounted 4 times.

Statistical mounts for A.B.C
= (mounts for volume 333333/# of data sets on volume 333333)
+ (mounts for volume 444444/# of data sets on volume 444444)
= 3/2 + 4/1 = 5.5 statistical mounts

This example shows that data sets can change their characteristics during the reporting period of the SMF sample. Since statistical mounts reflect all mounts for a data set during the time of the study, the number of statistical mounts for data set A.B.C is 5.5; the volume mount analyzer calculates mounts for both versions of A.B.C, first as multifile on volume 333333 and then as single file on volume 444444.

Statistical Volumes

A statistical volume is a calculation that shows the number of volumes directly associated with the current version of a data set. For this calculation, GFTAVMA divides a volume equally among the number of data sets or partial data sets residing on that volume.

Statistical volumes reflect only the volumes that contain the latest version of the data set. On the other hand, statistical mounts reflect all the mounts attributed to a data set during the reporting period. In this sense, the data set is a cumulative value. See “Statistical Mounts” on page 30 for information on statistical mounts.

Unique Volume Serial and Statistical Volumes

The total number of unique volume serials reported in GFTASRT2 often exceeds the number of statistical volumes reported in GFTASRT3. In GFTASRT2 (see Figure 17 on page 54), the number of UNIQUE VOLSERS is the count of all volumes used during the analysis period.

Figure 12 provides a good example. If you run GFTAVMA against one single-file, single-volume data set that was created five times on five different volumes, and if those volumes were never reused, the count of UNIQUE VOLSERS in GFTASRT2 would be five.

On day one of the study, data set D.E.F resided on volume 1.

VOLUME 1

D.E.F

 old version

On day four, a newer version of data set D.E.F resided on volume 2.

VOLUME 2

D.E.F

 old version

On day 12, another version of data set D.E.F resided on volume 3.

VOLUME 3

D.E.F

 old version

On day 23, an even newer data set D.E.F resided on volume 4.

VOLUME 4

D.E.F

 old version

On day 30, the latest version of data set D.E.F resides on volume 5.

VOLUME 5

D.E.F

 current version

Five different versions of data set D.E.F resided on five different volumes during the time of the study.

The GFTASRT2 report shows the total count of volumes on which all versions of a data set resided on during the period of the study (the number of unique volsers).

For data set D.E.F, UNIQUE VOLSERS = 5

Figure 12. Count of unique volsers for GFTASRT2

However, the statistical volume count in GFTASRT3 (see Figure 19 on page 58) would be one, which is the count of the statistical volumes needed to hold the latest version of the data set as shown in Figure 13 on page 33.

At the time of the study, the current version of data set D.E.F is the version residing on volume 5.

VOLUME 5 D.E.F current version

Volume 5 contains the current version of data set D.E.F, and GFTAVMA assumes that volumes 1 to 4 were returned to scratch because they contained old versions of data set D.E.F.

GFTASRT3 shows the count of the volumes that the current version of a data set resides on during the study (statistical volumes).

For data set D.E.F, STATISTICAL VOLUMES = 1

Figure 13. Count of statistical volumes for GFTASRT3

Zero Statistical Volumes

It is possible to have zero statistical volumes when a tape is reused. For example, on Monday, data set A resides on Volume 1 and causes one mount. On Wednesday, data set B resides on Volume 1 and causes one mount. Thus, there are two mounts for one volume. Since Volume 1 is reused, data set A has zero statistical volumes and one statistical mount, while data set B, the current version, has one statistical volume and one statistical mount.

Example of Calculating Statistical Volumes

VOLUME 111111 A.B.C

Statistical volumes for A.B.C = 1.0

Since data set A.B.C is a single-file, single-volume data set that only resided on volume 111111 during this analysis, the number of statistical volumes is 1.0.

Example for Multiple Versions of a Data Set

VOLUME 222222 A.B.C ' Day 1

VOLUME 333333 A.B.C " Day 2

Statistical volumes for A.B.C = 1.0

Data set A.B.C was written to Volume 222222 on Day 1 of the analysis. On Day 2, an updated version of data set A.B.C was written to a different volume (Volume 333333). Volume 222222 was then *erased-on-scratch* and made available for reuse. For the analysis, A.B.C occupied 1.0 statistical volume.

Example for Multifile and Multiple Version Data Sets

VOLUME 444444

X.Y.Z '

 Day 1

VOLUME 555555

X.Y.Z "	A.B.C
---------	-------

 Day 2

Statistical volumes for X.Y.Z = 0.5

Statistical volumes for A.B.C = 0.5

An earlier version of X.Y.Z resided on another volume (444444) before volume 444444 was erased-on-scratch. However, data sets A.B.C and X.Y.Z" currently share a single volume (555555).

Since the volume mount analyzer only calculates statistical volumes for the latest version of a data set, A.B.C occupies 0.5 statistical volumes, and X.Y.Z occupies 0.5 statistical volumes. Volume 444444 has been erased-on-scratch, since it no longer contains current data.

Open Counts

For each type of data set, the program that references it issues one open (one open count). However, a single-file, multivolume data set is likely to have multiple opens. This is because when processing reaches the end of the first volume, open/close/end-of-volume (O/C/EOV) issues another open to move to the next volume. So a multivolume data set can have two or more opens depending on the number of volumes on which it resides.

As the following example illustrates, the number of opens the volume mount analyzer reports for a single-file, single-volume data set does not equal the number of opens for a single-file, multivolume data set.

Open issued

VOLUME 111111

A.B.C

One open for each reference. Open count = 1

Open issued EOV open issued

VOLUME 111111

X.Y.Z (1)

 VOLUME 222222

X.Y.Z (2)

Two opens for each reference. Open count = 2

In this example, data set A.B.C. is opened once by the referencing program. It thus has an open count of one. Data set X.Y.Z is also opened once by the referencing program. But when the end of volume is reached, O/C/EOV issues a second open to move to volume 222222. As a result, data set X.Y.Z has an open count of two.

Tape Cartridge Usage

The volume mount analyzer does not know the physical length of a data set on a tape because the volume mount analyzer reads and analyzes SMF records, not the physical tape cartridges represented by the SMF records. Therefore, the volume mount analyzer data set size is the size of the data transfer to the tape cartridge.

The data set size calculation equals block count multiplied by block size. In most cases, the data transfer size equals the data set size. But since this calculation is an estimate, the two sizes are not always equal. If the data set is extended with `DISP=MOD`, the calculated size may be smaller than the actual size. If portions of the data set are reread or rewritten without closing and reopening the data set, the calculated size might be larger than the actual size. Examples are programs that reposition the tape with the `BSP`, `POINT`, `CNTRL`, `EXCP`, or `CLOSE TYPE=T` macros.

Since the type of tape cartridge, the compaction value, and the recording technique affect physical length of a data set on a tape, you need to be aware of the tape cartridge size, the compaction rate, and the recording technique.

To understand how data is placed on a tape, see *z/OS DFSMS: Using Magnetic Tapes*.

Volumes Without IDRC

Table 3 shows the capacities of tape volumes not compacted with improved data recording capability (IDRC).

Table 3. Capacities of Tape Volumes without IDRC Compaction

Type of Tape Volume	Size in Megabytes
3420 reel	169
3480 (348X)	208
3490E with 36 tracks (TAPELEN = 1)	416
3490E with 36 tracks and Enhanced Capacity Cartridge System Tape (TAPELEN = 2)	832

Cartridges With IDRC

Table 4 shows the capacities of tape cartridges compacted with improved data recording capability (IDRC), using a compaction factor of 3 (`IDRCFACTOR=3`).

Table 4. Capacities of Tape Cartridges with an IDRC Compaction Factor of 3

Type of Tape Cartridge	Size in Megabytes
348X (3490)	625
3490E with 36 tracks (TAPELEN = 1)	1 250
3490E with 36 tracks and Enhanced Capacity Cartridge System Tape (TAPELEN = 2)	2 500

Tape cartridge sizes range from 169MB for tape cartridges of the smallest capacity to 2500MB for those of the largest capacity. If you wish to take full advantage of the larger capacity tape cartridges and reduce your number of mounts, you need to use tape mount management and let DFSMSHsm place the data sets on tapes. Then you can fill these large capacity tape cartridges with data sets of a smaller size.

See *z/OS DFSMS: Implementing System-Managed Storage* for more information on how to implement advanced cartridge hardware.

Data Set Usage Types

GFTAXTR assigns a usage type to all data sets that the volume mount analyzer analyzes based on their naming convention, data set attributes, or patterns of read/write activity. There are six usage types: HSM, temporary, backup, single, BCOPY, and active.

There are also six categories of tape mount management data sets. See “Categories of Tape Mount Management Data Sets” on page 38 for more information. The tape mount management categories are different from data set usage types, but the usage types help determine the tape mount management categories in certain cases.

Order for Assigning Usage Types

GFTAXTR assigns usage types in the following order:

1. HSM
2. Temporary
3. Backup
4. Single
5. BCOPY
6. Active

You can filter on the usage types using the **USAGE** keyword. See “USAGE Keyword” on page 78 for more details.

HSM

If DFSMSHsm writes data sets to tape, GFTAXTR identifies them as USAGE(HSM) data sets. Because DFSMSHsm uses tape differently from other applications, any data sets written to tape by DFSMSHsm are included in USAGE(HSM).

DFSMSHsm Tape Usage

DFSMSHsm typically transfers large amounts of data. Therefore, GFTAXTR needs to separate these data sets. If the DFSMSHsm data sets are placed in active or backup categories, they can drastically influence any averages, totals, or summary information, which might cause the volume mount analyzer to draw inaccurate conclusions about those categories.

The most common usage of tape by DFSMSHsm is the single-file format (SFF) data sets used for migration level 2 or backup.

When migrating or backing up user data sets to tape with SFF, DFSMSHsm places data in blocks of 16384 bytes in these large, multivolume data sets. The SFF data sets are named hlq.HMIGTAPE.DATASET or hlq.BACKTAPE.DATASET. They are not cataloged and can span up to 255 tape volumes. DFSMSHsm continues writing user data sets in a single SFF data set until it is full. At that point, DFSMSHsm opens another SFF data set with the same name.

GFTAVMA Special Considerations for DFSMSHsm

The HSM usage pattern requires special consideration for GFTAVMA. For these data sets only, the cumulative gigabytes transferred represents the total amount of data written for output only over the entire sample studied. For all other data sets,

GFTAVMA uses the last quantity of data written as the current size of the data set. If there are no writes, the size of the data set is the amount of data transferred in the last open for output.

If migrations, recalls or backups, and recovery are the only activities occurring to the SFF data sets, GFTAVMA can measure how well DFSMShsm is using tape. However, if there is tape RECYCLE activity during the time studied, these opens, reads, and writes are added into the cumulative totals shown for the data set. Further, tape RECYCLE processes smaller amounts of data per read or write, which deflates the average megabytes/volume statistics.

DFSMShsm also uses various other utility and backup data sets, such as control data set backups, that are written to tape. Since DFSMShsm has special data set uses and sizes, GFTAXTR categorizes all DFSMShsm data sets as USAGE(HSM).

Temporary

Next, GFTAXTR assigns a usage type of *temporary* to data sets that have a naming convention matching the system-derived temporary name of SYS*.R*.*. Or it assigns the same data type to those data sets, regardless of their name, allocated with a disposition of (NEW,DELETE), indicating that they will be deleted at the end of the job. GFTAXTR categorizes these data sets as USAGE(TEMPORARY).

Backup

Then, GFTAXTR designates data sets that have been opened only once for output during the time period of the study as USAGE(BACKUP). Many of the USAGE(BACKUP) data sets are referenced on a single calendar day, but GFTAXTR separates them before the USAGE(SINGLE) data sets because the USAGE(BACKUP) data sets might migrate earlier, after 0 days, than USAGE(SINGLE) data sets.

USAGE(BACKUP) data sets are eligible to migrate to migration level 2 after 0 days on *primary storage* and 0 days on migration level 1, in other words, directly to migration level 2.

Single

After GFTAXTR has assigned the HSM, temporary, and backup usage types, it assigns USAGE(SINGLE).

GFTAXTR designates a data set as USAGE(SINGLE) if all references to it occur within a single calendar day. In doing so, GFTAXTR can identify data sets that can safely migrate to migration level 2 after one day because they are unlikely to be used again.

Unlike backup data, a USAGE(SINGLE) data set may have read activity. It does not matter whether the data set was opened for read or write, or how many opens were attributed to the data set. USAGE(SINGLE) requires, however, that the data is referenced only on a single calendar day.

USAGE(SINGLE) usually represents more than 50% of the data sets that can migrate directly to migration level 2 after one day.

BCOPY

GFTAXTR designates data sets with an access pattern of one open for input and one open for output as backup copy (BCOPY). Although these data sets do not

match the definition of backup data sets, in actuality, many BCOPY data sets are indeed a form of backup. This includes backups that have been recovered as well as backups that are written and then immediately copied to produce a vault copy.

If the open for input immediately follows the open for output on the same day, GFTAXTR categorizes the data set as USAGE(SINGLE). If the opens do not occur within the same day, GFTAXTR categorizes the data set as USAGE(BCOPY). This is a useful category to filter on to determine what this pattern of access means to your installation.

When calculating the values for the Estimate Report (see “The Estimate Report” on page 106), the volume mount analyzer categorizes BCOPY data sets as the tape mount management data set category of active and honors the **LEVEL0AGE** and **LEVEL1AGE** keyword specifications to determine where each data set should reside in the storage hierarchy.

Active

GFTAXTR categorizes data sets that do not fall into any of the above categories as USAGE(ACTIVE). These data sets have a mix of input and output activity and, typically, a higher number of mounts for a smaller amount of data than the data sets in other categories.

Do not directly access USAGE(ACTIVE) data sets from tape unless they are very large. These data sets should be part of the storage hierarchy, managed by DFSMS, allowed to automatically migrate to tape, and based on their age and pattern of usage.

You can use the Management Class Report (see “The Management Class Report” on page 125), to determine the best **LEVEL0AGE** and **LEVEL1AGE** values for any given set of data based on the actual reference pattern of the data sets.

Categories of Tape Mount Management Data Sets

In the tape mount management methodology, GFTAVMA categorizes all data sets according to their usage patterns. Data sets are categorized into one of six possible tape mount management categories.

In a volume mount analyzer study, tape mount management categories affect the output of the Estimate, IDRC, and Management Class Reports. See “Interpreting the Simulation Reports” on page 106 for more details.

Order for Assigning Tape Mount Management Data Set Categories

When calculating the simulation reports, GFTAVMA divides all tape data sets into six categories. GFTAVMA assigns the tape mount management categories in the following order:

Exceptions

Categorized using **EXCLUDE** filters and the **SPLIT** keyword. See “Include/Exclude Filters” on page 44 and “SPLIT Keyword” on page 151 for more information.

HSM Categorized because the data sets were created by DFSMSHsm and have the USAGE(HSM) attribute.

LARGE

Categorized using the **LARGE** keyword. All data sets whose sizes are greater than or equal to the value of the **LARGE** keyword are large. See “LARGE Keyword” on page 144 for more information.

Temporary

Categorized because the data sets have a temporary data set name or were created with `DISP=(NEW,DELETE)`.

Backup

Categorized by the **CLASSIFY** keyword or because the data sets were built by a backup program and categorized as `USAGE(BACKUP)`. See “CLASSIFY Keyword” on page 135 for more information.

Active Categorized by the **CLASSIFY** keyword or because these data sets have not been assigned other tape mount management data set categories. See “CLASSIFY Keyword” on page 135 for more information.

Tape Mount Management Data Set Categories in the Estimate Report

The Estimate Report accounts for data sets based on the categories of active or backup and simulates the management of data sets based on their particular tape mount management data set categories. See “The Estimate Report” on page 106 for more information.

Active Data sets are *modelled* as migrating through the storage hierarchy once every 24 hours according to the management class keywords **LEVEL0AGE** and **LEVEL1AGE**.

Data sets categorized as `USAGE(ACTIVE)` or `USAGE(SINGLE)` are modelled as the tape mount management category of active for the Estimate Report.

Backup

Data sets are modelled as migrating once every 24 hours directly to migration level 2 from primary DASD after zero days of residency. The Estimate Report assumes that the data sets migrate every 24 hours, not hourly.

Data sets categorized as `USAGE(BACKUP)` are modelled as the tape mount management category of backup for the Estimate Report.

LARGE

Data sets whose sizes are greater than or equal to the **LARGE** keyword value are modelled as bypassing the tape mount management DASD buffer and going directly to tape.

Any data sets larger than the **LARGE** value are modelled as the tape mount management category of large for the Estimate Report. Data sets from any usage type can be large depending on their size.

HSM Data sets created by `DFSMSHsm` are modelled as bypassing the tape mount management DASD buffer and going directly to tape.

Data sets categorized as `USAGE(HSM)` are modelled as the tape mount management category of HSM for the Estimate Report.

Temporary

Data sets are modelled as going directly to the tape mount management DASD buffer for the life of the job and then expiring.

Data sets categorized as USAGE(TEMP) are modelled as the tape mount management category of temporary for the Estimate Report.

Exceptions

Data sets that remain on tape.

Categorization of Data Sets

GFTAVMA assigns tape mount management data set categories based on the data set name, data set size, the pattern of read and write activity to a data set, and the keywords you have specified. See Figure 14 to understand the order of assignment.

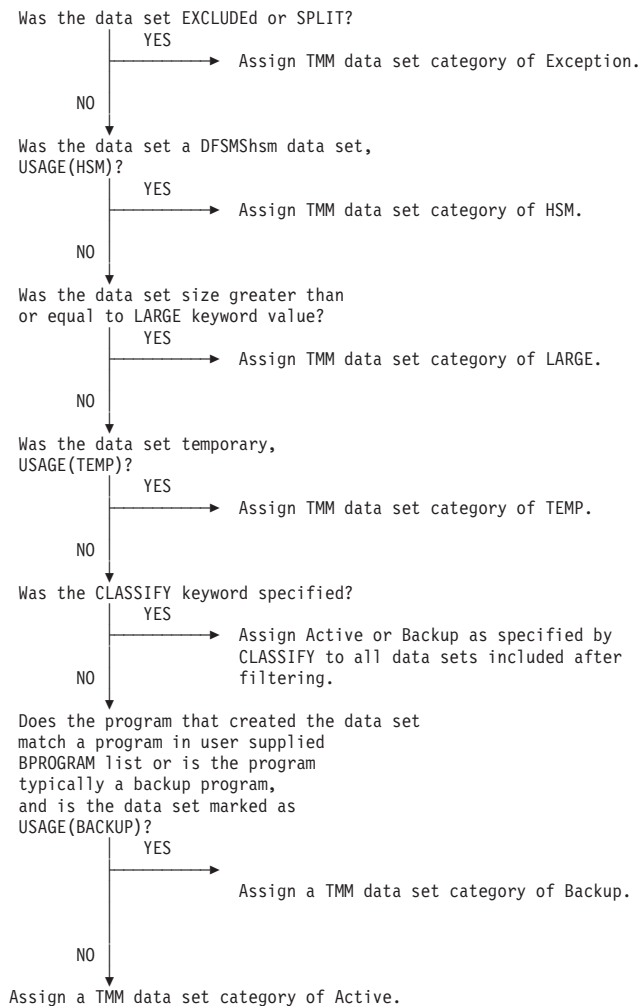


Figure 14. Order of Assigning Tape Mount Management Categories

Chapter 4. Analyzing Your Tape Usage with GFTAVMA

This chapter helps you understand the GFTAVMA program and explains how to analyze your tape usage with GFTAVMA.

For a set of SMF data, you run GFTAXTR only once to create the extract file. However, you run GFTAVMA many times to determine which applications would get the most benefit from tape mount management for the lowest cost. These runs also generate the list of ACS filters you can use to capture data sets and to determine the appropriate management class and data class for each application.

Understanding the GFTAVMA Program

GFTAVMA is the program name of the volume mount analyzer. For input, GFTAVMA requires GFTAXTR's output data set, which is defined to GFTAVMA by the XTRIN DD statement in GFTAVMA's JCL.

Note: You must pass the GFTAXTR output data set to GFTAVMA unaltered. Do not attempt to edit the GFTAXTR data set. You cannot concatenate or combine multiple GFTAXTR data sets as input to GFTAVMA.

Requirements for Using GFTAVMA

To run GFTAVMA, you must have a z/OS operating system.

You must also understand how to create useful filters and how to use the optional keywords for GFTAVMA. See "Filtering Your Input with GFTAVMA Keywords" on page 44 for an explanation of creating filters.

See Chapter 5, "Using GFTAVMA Keywords to Filter Your Input," on page 63 and Chapter 7, "Using GFTAVMA Keywords to Tailor Your Output," on page 131 for more information on keywords.

GFTAVMA Output Reports

GFTAVMA generates the following output:

- Summary reports that GFTAVMA automatically creates every time you run the volume mount analyzer
- Optional detailed reports that GFTAVMA creates depending on the report keywords that you specify for a run

See "REPORT Keyword" on page 148 for more information on the report keywords.

GFTAVMA generates the summary reports and the optional detailed reports during the three phases of GFTAVMA processing. If you request the detailed reports using the **REPORT** keyword, you get one output data set containing the three summary reports and any optional detailed reports you requested.

You can use GFTAVMA reports to determine which data sets are good candidates for tape mount management.

Because you can generate many unique GFTAVMA reports during your analysis, create DASD data sets to manage them more easily.

Understanding GFTAVMA Processing Phases

There are three GFTAVMA processing phases. The reports generated during each phase appear in the output data set in the order GFTAVMA generates them:

- During phase 1, GFTAVMA generates the GFTASRT1 summary report and, if requested, the Usage Report.
- During phase 2, GFTAVMA generates the GFTASRT2 summary report and, if requested, the Volume Report.
- During phase 3, GFTAVMA generates the GFTASRT3 summary report and, if requested, any or all other optional reports.

If you request the Volume Report with the **REPORT** keyword, the order of the reports is GFTASRT1, GFTASRT2, the Volume Report, and GFTASRT3. See “Interpreting GFTAVMA Summary Reports” on page 51 and Chapter 6, “Producing Detailed Analysis Reports,” on page 81 for more information.

Figure 15 shows the filtering and processing steps of GFTAVMA.

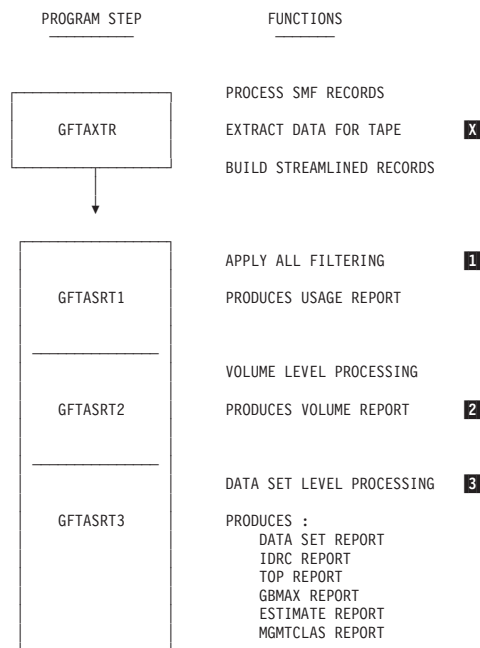


Figure 15. The Three Phases of GFTAVMA Processing

X GFTAXTR generates the extract output data set defined by the XTRCIN DD statement in the GFTAXTR JCL. This data set must be passed to GFTAVMA in the exact form as it was created. If a data set is created by removing a subset of the GFTAXTR file or by merging several GFTAXTR files, GFTAVMA terminates with an error.

1 Each of the GFTAVMA processing steps generates output. See “Interpreting GFTAVMA Summary Reports” on page 51 for more information about the summary reports.

All filtering based on keywords and sub-keywords, such as **INCLUDE**, **EXCLUDE**, **MAXSIZE**, and **USAGE**, is done at the beginning of the GFTASRT1 phase. As a result, GFTAVMA generates all reports with filtered data. The Usage Report, for example, shows only the number of devices allocated to those tape volumes still included in the analysis after GFTAVMA has applied filtering.

You control the filtering of the analysis reports using keywords on a control statement in the JCL member, GFTAVMAP. You can submit GFTAVMAP multiple times using different keywords for each run.

See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 and Chapter 7, “Using GFTAVMA Keywords to Tailor Your Output,” on page 131 for the rules and syntax of the optional keywords.

2 Volume level processing examines the data sets from the volume perspective. The Volume Report shows the amount of data that was processed to each volume. If a volume was used more than once, the last data written to the volume is considered to be what is on the volume.

The intention of the Volume Report is to show how full the volumes are. Therefore, the numbers on the volume reports and the data set reports do not match.

3 Data set level processing examines the data sets from the data set perspective. This level of processing considers the most *current* version of each data set to be the last one written during the time period. If the data set was not created during the time period of the study, then GFTAVMA uses the last read as the most current version of the data set.

GFTAVMA accumulates mounts and opens for all versions analyzed in the sample, but it considers size and volume residency only for the current version.

Preparing to Run GFTAVMA

Before you run GFTAVMA, look at the volume mount analyzer messages in the GFTAXTR output report defined by the SYSPRINT DD statement in the GFTAXTR JCL. Also check the following in SYSPRINT:

- Do you have type 14 and 15 tape records?
- What are the start and end dates for the GFTAXTR run?
- How many records were processed?
- Are all dropped records accounted for?
- How many SMF records are in the GFTAXTR output file?
- How many tape jobs were processed?
- How many input and output records are there?
- What systems are associated with the sample?
- What are the error and informational messages?

This information ensures that you collected a representative sample of System Management Facility (SMF) data for input to GFTAVMA. See Appendix D, “Volume Mount Analyzer Messages,” on page 167 for more information.

z/OS DFSMS: Implementing System-Managed Storage explains the most useful order for requesting the detailed reports. Read about tape mount management in this book before you use GFTAVMA for the first time.

Filtering Your Input with GFTAVMA Keywords

The volume mount analyzer program, GFTAVMA, provides a comprehensive set of *filters* for you to study your tape data set activity and apply the tape mount management methodology. Filters are the criteria that data sets must meet in order to be included in your analysis.

You select optional keywords at GFTAVMA execution time to filter the input to GFTAVMA. See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 for more information about the filtering keyword rules and syntax.

Filter Types

The volume mount analyzer has two types of filters:

- limiting
- include/exclude

Limiting Filters

All data is first filtered by the limiting filters, which restrict the data in your analysis. GFTAVMA invokes these filters in the following order:

1. MAXSIZE
2. MINSIZE
3. MOUNT
4. FILE
5. USAGE
6. DATE
7. TIME

The limiting filters act as a funnel. For example, **MAXSIZE** filters the complete set of input data. **MINSIZE** filters the subset of data that remains following **MAXSIZE** filtering. **MOUNT** filters the subset of data that remains following **MINSIZE** filtering and so on.

If you specify the following keywords:

```
MAXSIZE(99) USAGE(NONHSM) DATE(01/01/04,01/02/04)
```

your input data is limited to non-HSM data sets that are less than 100MB and were created during the first two days of January. Each of these limiting filters further restricts the data sets that are candidates for selection. In this way, the size of the SMF sample gets smaller and smaller.

Include/Exclude Filters

The second type of filter is the include/exclude filter. These filters affect only the subset of data that remains following the limiting filters. They are slightly more complex than the limiting filters. They use **INCLUDE** and **EXCLUDE** sub-keywords so that you can be more selective when filtering your data. Further, these filters use Boolean “OR” logic to select data sets from the smaller sample produced by the limiting filters.

Categories of Include/Exclude Filters

The volume mount analyzer program has two categories of include/exclude filters:

- primary
- secondary

Primary Filters

Primary filters are the filters you can use for ACS routines and for tape mount management. A data set meeting the criteria defined by any of these primary filters is included in or excluded from the analysis sample.

The five primary filters, when used together, interact using Boolean "OR" logic:

- ACCOUNT
- DATASET
- EXPDT
- JOBNAME
- PROGRAM

During GFTAXTR processing, the latest creating **JOBNAME**, **PROGRAM**, **ACCOUNT**, and data set **EXPDT** are saved in all versions of the data set record if the data set was actually created during the time period represented by the SMF sample. Otherwise, the **JOBNAME**, **PROGRAM**, **ACCOUNT**, and data set **EXPDT** of the first reference are saved in all versions of the data set record. Therefore, when you filter your data using these primary filters, you filter in or out all versions of the data sets based on the creating **JOBNAME**, **PROGRAM**, **ACCOUNT**, and data set **EXPDT**.

Secondary Filters

Data sets that have not been excluded by limiting or primary filtering are further filtered by the secondary filters. You can use secondary filters to select specific units, volumes, or systems. These filters help you reduce your output.

The four secondary filters, when used together, interact using Boolean "OR" logic:

- SYSTEMID
- UNIT
- UNITADDR
- VOLUME

Example of Primary and Secondary Filters

When you use primary and secondary filters together, the data is filtered by the primary **DATASET** filter first, and then by the secondary **UNIT** filter.

```
DATASET (INCLUDE (AAA.**)) UNIT (INCLUDE (3490E))
```

This combination would include data sets with a high level qualifier of AAA. Of this group of data sets, it would include only those referenced on a 3490E tape drive.

Note: As this example shows, when you use a primary and a secondary filter together, they do not interact with "OR" logic. Primary filters interact using "OR" logic with other primary filters, and secondary filters interact using "OR" logic with other secondary filters. However, primary filters interact with secondary filters using "AND" logic.

Include/Exclude Filter Syntax

You can use either or both of the **INCLUDE** and **EXCLUDE** sub-keywords on the include/exclude filtering keywords. You can specify filter strings on the **INCLUDE** and **EXCLUDE** sub-keywords. The items in the filter strings must be separated by commas. For example:

```
DATASET (INCLUDE (ABC*), EXCLUDE (ABC001,ABC002))
```

This example includes any data sets beginning with the characters **ABC** except **ABC001** or **ABC002**.

Note: Do not put quotation marks around sub-keywords in your include-exclude filters because the volume mount analyzer assumes the quotation marks are part of the sub-keywords.

Global Data Set Name Characters in Filters

You can use these global data set name characters in your filters:

Single asterisk (*)

represents one of the following:

- A single data set name qualifier (**A.*.B**)
- Zero to 8 characters within a data set name qualifier (**ABC***)
- Zero to 8 characters within a job name or program name
- Zero to 6 character within a volume serial

Double asterisk (**)

represents zero to 22 qualifiers within a qualified data set name (**A.**.B**)

Percent Sign (%)

represents one, and only one, character

Examples of Global Data Set Name Characters

**LOAD	Any qualified name where the last qualifier is LOAD .
ABCD.*	Any name with only 2 qualifiers where the first qualifier is ABCD . Any name starting with ABCD that has more than 2 qualifiers will be ignored.
SYS1.**	Any qualified name where the first qualifier is SYS1 .
XYZ**.*LIST	Any qualified name where the first 3 characters of the first qualifier are XYZ , and the last 4 characters of the last qualifier are LIST with any number of qualifiers between them.
348*	3480 , 348X
SYS1**.*ASM	Any name with 4 qualifiers where the first qualifier is SYS1 , and the last qualifier is ASM .
**B%%.G*V00	Any <i>generation data set (GDS)</i> where the qualifier immediately preceding the generation suffix consists of a 'B' followed by any two other characters.

Examples of Include/Exclude Filters

If you specify the **DATASET** keyword as follows:

```
DATASET (INCLUDE (AAA.B%B.**), EXCLUDE (AAA.BZB.**))
```

only data sets with a high-level qualifier of **AAA** and a second-level qualifier beginning and ending with **B** would fit the **INCLUDE** pattern. From this group of selected data sets, **GFTAVMA** would **EXCLUDE** those whose second-level qualifier is **BZB**.

If you specify a **DATASET** keyword with no **INCLUDE**s, only **EXCLUDE**s, GFTAVMA compares all the data that was not previously filtered against the **EXCLUDE** list and excludes from the analysis the data sets that fit the **EXCLUDE** pattern.

The following example shows the **DATASET** keyword with an **INCLUDE** that works, but an **EXCLUDE** that is superfluous:

```
DATASET (INCLUDE (AAA.B%B.**), EXCLUDE (AAA.CZC.**))
```

Only data sets with a high-level qualifier of AAA and a second-level qualifier beginning and ending with B fit the **INCLUDE** pattern. However, the **EXCLUDE** pattern says exclude data sets with a high-level qualifier of AAA and a second-level qualifier of CZC. Unlike the previous example, this pattern is not a subset of the **INCLUDE** pattern. Thus, the **EXCLUDE** sub-keyword is superfluous because these data sets would be excluded anyway since they do not match the include pattern.

If you specify two primary include/exclude filters together:

```
DATASET (INCLUDE (AAA.**)) JOBNAME (INCLUDE (FQX73B))
```

they interact using Boolean "OR" logic. Therefore, the results of this combination would include all data sets with a high-level qualifier of AAA and all data sets created by the job named FQX73B.

Example of Generation Data Set Filters

To include GDSs, specify

```
DATASET (INCLUDE (AAA.BBB.CCC.G%%V%%))
```

This example includes GDSs from the *generation data group* (GDG) named AAA.BBB.CCC. Each GDS, within the GDG, has a suffix that identifies the absolute generation and version. The generation and version are in the form GxxxxVyy, where xxxx is an unsigned 4-digit decimal generation number (0001 through 9999) and yy is an unsigned 2-digit decimal version number (00 through 99).

GFTAVMA does not filter generation data sets by relative generation number, such as AAA.BBB.CCC(-1) or AAA.BBB.CCC(0). Additionally, if you specify the GDG name, AAA.BBB.CCC, without the G%%V%% suffix, GFTAVMA will not know it as a GDG name. Consequently, you can't specify the GDG name and expect GFTAVMA to include or exclude all of the GDSs within the GDG.

Note: Do not put quotation marks around sub-keywords in your include/exclude filters because the volume mount analyzer assumes the quotation marks are part of the sub-keywords.

Example of Multiple Filters

If you specify

```
DATASET (INCLUDE (A.B.C))  
PROGRAM (INCLUDE (ADRDSSU))  
UNIT (INCLUDE (3490E))
```

this set of keywords will include all records that have a data set name of A.B.C and all records created by the program named ADRDSSU. Of the set of data that matches either of these two keywords, only those with unit type of 3490E are included.

In this example, the primary filters, **DATASET** and **PROGRAM**, interact with each other using Boolean "OR" logic. Then the two primary filters interact with the secondary filter, **UNIT**, using Boolean "AND" logic.

Filter Cautions

Primary and secondary filters sometimes do not mix well when you use them together.

For example, if you combine the primary **DATASET** filter and the secondary **UNITADDR** filter in the following way, you might get inaccurate results.

```
DATASET (INCLUDE (AAA.**)) UNITADDR (INCLUDE (0580))
```

The results of this combination would include data sets with the high-level qualifier of AAA, but of this group of data sets, only those on the device with the address 0580 would be included. This might cause you to artificially exclude some AAA data sets from the analysis, which would affect your report results.

Submitting the GFTAVMA JCL

After you have determined the filtering and report keywords you want to use for your volume mount analyzer run, you must submit the GFTAVMA job. To submit the GFTAVMA job, you:

- Copy the JCL member, GFTAVMAP, from SYS1.SAMPLIB
- Change *only* the indicated fields
- Submit the GFTAVMA JCL for execution using the TSO SUBMIT command

Note: It is *crucial* to change only the indicated fields. If you use this JCL as is, you ensure the space calculation is accurate. GFTAVMA uses system determined block size and average record allocation. Since SORT is invoked multiple times, it is important not to have RLSE on the SORTxxxx and XTRCxxxx files.

For more information on using JCL, see *z/OS MVS JCL Reference*.


```

//GFTAVMA JOB (ACCT#,'ACCOUNTING-INFORMATION'),'PROGRAMMER NAME', 1
//      MSGCLASS=A,MSGLEVEL=(1,1),USER=&SYSUID,NOTIFY=&SYSUID
//*
//GFTAVMA PROC HLQ=,          * GFTAXTR FILE HIGH LEVEL QUALIFIER
//      RUN=,                * GFTAXTR FILE LOW LEVEL QUALIFIER
//      REGVMA=6144K,        * DEFAULT REGION FOR 'GFTAVMA'
//      TAPES='20000,3000', * DEFAULT NUMBER OF TAPE MOUNTS
//      UNIT=SYSDA,         * DEFAULT ESOTERIC UNIT NAME
//      FILTERS=1000        * MAXIMUM NUMBER OF FILTERS
//*****
//*          GFTAVMA -- VOLUME MOUNT ANALYSIS PROGRAM          *
//*          *          *          *          *          *          *
//*****
//* THIS JCL IS USED TO RUN THE 'GFTAVMA' ANALYSIS PROGRAM. THIS *
//* PROGRAM HAS THE CAPABILITY TO ANALYZE TAPE USAGE DATA FROM SMF. *
//*          *          *          *          *          *          *
//* 'GFTAVMA' USES THE INPUT FILE PRODUCED BY THE 'GFTAXTR' PROGRAM. *
//* 'GFTAVMA' WILL TERMINATE PROCESSING IF NO RECORDS WERE FOUND IN *
//* THE INPUT FILE (I.E., THE 'XTRIN' DD) FROM THE 'GFTAXTR' UTILITY. *
//*          *          *          *          *          *          *
//*****
//*          *          *          *          *          *          *
//VMA EXEC PGM=GFTAVMA,REGION=&REGVMA,PARM='FILTR#(&FILTERS)' 2
//SORTIN DD UNIT=&UNIT,
//      SPACE=(1000,(&TAPES)),AVGREC=U
//SORTOUT DD UNIT=&UNIT,
//      SPACE=(1000,(&TAPES)),AVGREC=U 3
//SORTWK01 DD UNIT=&UNIT,
//      SPACE=(350,(&TAPES),,CONTIG,ROUND),AVGREC=U
//SORTWK02 DD UNIT=&UNIT,
//      SPACE=(350,(&TAPES),,CONTIG,ROUND),AVGREC=U
//SORTWK03 DD UNIT=&UNIT,
//      SPACE=(350,(&TAPES),,CONTIG,ROUND),AVGREC=U
//SYSOUT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=*
//VMACHART DD SYSOUT=* 4
//VMAINCL DD DUMMY,DCB=(DSORG=PS,RECFM=VB)
//VMAEXCL DD DUMMY,DCB=(DSORG=PS,RECFM=VB) 5
//VMACNTL DD DUMMY
//VMAFLTRS DD UNIT=&UNIT,
//      SPACE=(10,(&TAPES),RLSE),AVGREC=U
//XTRCIN DD UNIT=&UNIT,
//      SPACE=(1000,(&TAPES)),AVGREC=U
//XTRCOUT DD UNIT=&UNIT,
//      SPACE=(1000,(&TAPES)),AVGREC=U
//XTRCWK01 DD UNIT=&UNIT,
//      SPACE=(350,(&TAPES),,CONTIG,ROUND),AVGREC=U
//XTRCWK02 DD UNIT=&UNIT,
//      SPACE=(350,(&TAPES),,CONTIG,ROUND),AVGREC=U
//XTRCWK03 DD UNIT=&UNIT,
//      SPACE=(350,(&TAPES),,CONTIG,ROUND),AVGREC=U
//XTRIN DD DISP=SHR,DSN=&HLQ..GFTAXTR.&RUN

```

1 Substitute your own JOB statement.

2 During initialization, GFTAVMA obtains enough memory to store up to 1000 user-specified filters. If you need to use more than 1000 filters during GFTAVMA processing, you can set the value of the FILTR#(&FILTERS) parameter on the EXEC statement in **6**. The FILTERS parameter allows you to specify the number of filters that you want to supply as input to filter your SMF data. You can change the FILTR# value by changing the FILTERS specification. The valid range is 100 to 21474836.

If you remove the PARM='FILTR#(&FILTERS)' from the VMA EXEC statement, GFTAVMA defaults to 10,000.

3 All SORT/WORK files defined by the SORTxxx and XTRCxxx prefixed DD statements are sized according to the values supplied on the TAPES parameter. During GFTAVMA processing, DFSORT does not release space automatically on each call. Therefore, do not use the RLSE sub-parameter of the DD SPACE parameter for these JCL statements. See the note in the comments block regarding setting the TAPES parameter. These values supply the space quantities for all files required during the GFTAVMA run.

4 The **CHART** keyword is optional. When specified, the **CHART** keyword causes the GFTAVMA program to create a sequential data set referenced by the VMACHART DD statement.

In the GFTAVMAP JCL member, the VMACHART DD is allocated to a SYSOUT file. To override this specification, specify the **CHART** keyword after **7** and insert the following DD statement after the EXEC statement in **6** :

```
//VMA.VMACHART DD DSN=your.data.set,DISP=(NEW,CATLG),  
//   DCB=(LRECL=80,RECFM=FB),SPACE=(TRK,(1,1))
```

This statement directs the **CHART** output to a sequential data set that you can use to create charts using the graphics package of your choice.

If you do not request **CHART**, GFTAVMA does not put any output in the SYSPRINT file.

See “CHART Keyword” on page 134 for more information about the **CHART** keyword. See Appendix C, “Example of Chart Output,” on page 161 for examples of the output that **CHART** generates.

5 The **SPLIT** keyword is optional. Use **SPLIT** to filter out a large list of excluded data sets once rather than filtering them out each time.

The **SPLIT** keyword creates two optional sequential data sets. These data sets are identified by the VMAINCL and VMAEXCL DD statements in the GFTAVMA JCL. In the GFTAVMAP JCL member, GFTAVMA PROC has assigned these DD statements as dummy.

To use **SPLIT**, you must override these dummy assignments. To override the current GFTAVMA DD assignments, use the following DD statements to redirect them to a direct access or tape data set.

Place the DD statements for **SPLIT** after the EXEC GFTAVMA statement (**6**) in the sample JCL, and specify the **SPLIT** keyword after the VMA.VMACNTL DD statement (**7**).

```
//VMA.VMAINCL DD DSN=your.sequential.data.set.for.includes,  
//   DCB=(LRECL=316,RECFM=VB,DSORG=PS),  
//   DISP=(NEW,CATLG),  
//   SPACE=(.....),  
//   UNIT=...  
//*  
//VMA.VMAEXCL DD DSN=your.sequential.data.set.for.excludes,  
//   DCB=(LRECL=316,RECFM=VB,DSORG=PS),  
//   DISP=(NEW,CATLG),  
//   SPACE=(.....),  
//   UNIT=...
```

For subsequent GFTAVMA runs, update the XTRIN DD statement to point to either the included or excluded data set.

See “SPLIT Keyword” on page 151 for more information.

```

//*****
//*
//* YOU MUST SET "?HLQ" AND "?RUN" PROC PARAMETERS TO SPECIFY THE
//* INPUT DATA SET FROM THE 'GFTAXTR' RUN
//*
//* THE FOLLOWING DEFAULT VALUES ARE IN THE "GFTAVMA" PROC:
//*
//* REGVMA=6144K * DEFAULT REGION FOR 'GFTAVMA'
//* TAPES='20000,3000' * DEFAULT NUMBER OF TAPE MOUNTS
//* UNIT=SYSDA * DEFAULT ESOTERIC UNIT NAME
//* FILTERS=1000 * MAXIMUM NUMBER OF FILTERS
//*
//* NOTES:
//* 1) TO SET THE 'TAPES' VALUE PROPERLY FOR CORRECT SPACE
//* ESTIMATES ON ALL DD STATEMENTS IN THIS PROC, USE THE
//* NUMBER OF TYPE 21 RECORDS IN THE INPUT SMF SAMPLE AS
//* THE VALUE. THE 2ND VALUE SHOULD BE SPECIFIED AS A
//* SAFETY VALVE FOR OVERFLOW. FOR EXAMPLE, IF THE INPUT
//* SAMPLE HAD 28,653 TYPE 21 RECORDS, ONE MIGHT SET THE
//* TAPES='30000,5000'
//*
//* 2) THE 'FILTERS' VALUE WILL ALLOW THE GFTAVMA PROGRAM TO
//* ACQUIRE STORAGE FOR THE FILTER KEYS. YOU SHOULD
//* ESTIMATE THE MAXIMUM NUMBER OF INCLUDE/EXCLUDE FILTER
//* KEYS THAT YOU WILL SPECIFY IN THE GFTAVMA ANALYSIS RUN.
//*
//*****
//*
//* PEND
// EXEC GFTAVMA,HLQ=?HLQ,RUN=?RUN
//*
//*****
//* PUT YOUR 'GFTAVMA' KEYWORDS AFTER THE "VMACNTL" DD STATEMENT
//* BELOW.
//*
//*****
//*
//VMA.VMACNTL DD *

```

7

6 Specify the name of your GFTAXTR output data set, XTRCIN.

7 Specify keywords and filters here for the GFTAVMA run. These keywords are all optional. If you do not specify keywords or filters in the JCL, GFTAVMA analyzes all the input from GFTAXTR and only generates the summary reports, which are created for every run.

Note: *z/OS DFSMS: Implementing System-Managed Storage* describes the recommended order for requesting the detailed reports. Read these recommendations before you run GFTAVMA.

Troubleshooting When Obtaining GFTAVMA Output

If you did not get output from GFTAVMA, you probably used too many filters. That is, you over-filtered the GFTAXTR output. If you filtered out all SMF records, GFTAVMA had nothing to analyze. Check your list of filters to see what filters you need to change to include SMF records in the current run. See “Filtering Your Input with GFTAVMA Keywords” on page 44 for more information about filtering.

Look at the volume mount analyzer messages to determine why you did not get output. See Appendix D, “Volume Mount Analyzer Messages,” on page 167 for explanations of the messages.

Also check the levels for the GFTAXTR program and the GFTAVMA program. These levels must match. If they do not match, the volume mount analyzer terminates with an error message and does not produce output.

Interpreting GFTAVMA Summary Reports

Each major processing step of GFTAVMA generates its own output consisting of processing statistics, messages, tables, and summary information about the data.

Summary reports are automatically generated for every GFTAVMA run. GFTAVMA puts the summary reports in the data set defined by the SYSPRINT DD statement in the GFTAVMAP JCL. Normally, this data set is a SYSOUT data set.

Examples of the output summary reports from each processing step of GFTAVMA are shown in the following sections. These examples are generated by running the GFTAVMAP JCL member without specifying any optional report keywords. The actual content of the summary reports varies depending on the keywords used and the amount and type of data being processed. See “Submitting the GFTAVMA JCL” on page 48 for the JCL that generated these examples.

If you specify the **REPORT** keyword, in addition to the summary reports, you also get the detailed reports you requested. See Chapter 6, “Producing Detailed Analysis Reports,” on page 81 for more information.

All volume mount analyzer messages appear in the reports. See Appendix D, “Volume Mount Analyzer Messages,” on page 167 for more information.

Input Analysis Phase—GFTASRT1

This phase analyzes the input that GFTAVMA receives from GFTAXTR by building and filtering the SMF data. GFTAVMA includes or excludes all records based on the filters you use. See “Filtering Your Input with GFTAVMA Keywords” on page 44 for more information about filtering.

Note: To be sure that the GFTAVMA and GFTAXTR runs are for the same set of data, compare the times and system IDs on the GFTAXTR report to those on the summary reports. They must be the same.

GFTASRT1—Input Analysis Report

As shown in Figure 16 on page 53, GFTASRT1 is the part of the GFTAVMA summary report that shows the times and dates of the SMF input data sample and the GFTAXTR run. It shows the input records that GFTAXTR generated and the output records that GFTAVMA used when the filtering phase ended.

In this example report, keywords were not specified.

```

SMF INPUT SAMPLE: 33 DAYS
                  START: 05/27/2003 -- 12:02:01 A.M.
                  END:   06/28/2003 -- 11:59:18 P.M.

GFTAXTR RUN:
  START: 05/21/2004 -- 02:33:02 P.M. 1
  END:   05/21/2004 -- 02:57:22 P.M.

JES SYSTEM ID(S):

  SYSTEM ID  OPER SYS
  -----
  SYSA      MVS/ESA

GFTASRT1 -- INTERNAL RECORD BUILDING/FILTERING PHASE -- 05/21/2004 -- 03:44:29 P.M.

INPUT RECORD STATISTICS: (GENERATED BY THE 'GFTAXTR' PROGRAM)
  1 TYPE 00 RECORDS READ ('GFTAXTR' VERSION RECORD)
  1 TYPE 01 RECORDS READ (JES SYSTEM ID RECORD)
 17511 TYPE 14 RECORDS READ (OPEN FOR INPUT)
 67243 TYPE 15 RECORDS READ (OPEN FOR OUTPUT)
 38586 TYPE 21 RECORDS READ (TAPE DEMOUNTS)
-----
123342 TYPE XX RECORDS READ (TOTAL FROM 'XTRIN' FILE) 2

OUTPUT RECORD STATISTICS: (USED INTERNALLY BY THE 'GFTAVMA' PROGRAM)
 17511 TAPE DATA SET RECORDS (OPEN INPUT)
 67243 TAPE DATA SET RECORDS (OPEN OUTPUT)
 38586 TAPE VOLUME RECORDS
-----
123340 TOTAL OUTPUT RECORDS ('XTRCIN' FILE) 3

DATA SET/VOLUME/MOUNT/GB EXCLUSION STATISTICS:

DATA SET(S)  PERCENT  GB(S)  PERCENT  STAT  MNT(S)  PERCENT  STAT  VOL(S)  PERCENT
0            0.0%    0.0    0.0%    0.0    0.0%    0.0    0.0%    0.0    0.0%  TOTAL EXCLUDED 4

```

Figure 16. GFTASRT1 Summary Report—Input Analysis Phase

1 Be sure the date, time, and system ID on this report match the date, time, and system ID on the GFTAXTR report. See “GFTAXTR Output Report—Final Record Output Phase” on page 26.

2 These are the records built by the GFTAXTR program.

3 There is a difference of two between this total and the total in **2**. The two records deducted are the header records that contain the GFTAXTR version information, SMF record counts, system IDs, and time/date information. The input phase uses these two records but does not pass them on to the processing phases that follow.

4 Any data sets excluded because of filtering, such as **INCLUDE**, **EXCLUDE**, **USAGE**, or **MAXSIZE**, are listed here. Since no filters were specified for this run, nothing was excluded.

Volume Analysis Phase—GFTASRT2

During the volume analysis phase, GFTAVMA analyzes the data from the volume perspective and generates the GFTASRT2 summary report.

GFTASRT2—Tape Volume Statistics Report

GFTASRT2 is the part of the GFTAVMA summary report that shows how much of your tape media actually contain data. This report reflects the latest contents of the volume. The data sets that reside on a particular volume might be old versions of the data sets, but they are the current contents of the volume.

GFTASRT2 Report

In the volume analysis phase, only volume level processing is completed. Since the SMF sample represents a time period, it is possible for the same data set to have been re-written several times on multiple volumes. See "Statistical Volumes" on page 31 for an explanation of unique volsers.

For example, if data set 'X' is written ten times onto ten different tape volumes during the sample, all these tape volumes and their associated mounts are included in the volume level processing. GFTAVMA analyzes all of these volumes to show an accurate picture of how well your volumes are used over time, even though only one of the volumes contains valid data in this case.

Figure 17 shows the tape volume serial (volser) statistics. In this report, keywords were not specified.

```
GFTASRT2 -- TAPE VOLSER PROCESSING PHASE -- 05/21/1992 -- 03:44:30 P.M.
```

```
*****
*
*                                     GFTASRT2 UNIQUE TAPE VOLSER STATISTICS
*
*****
```

DEVICE TYPE	1 UNIQUE VOLSERS	2 # OF SF VOLSERS	3 # OF MF VOLSERS	RETURNED SCRATCH	STANDARD LABEL	NON-STD LABEL	ISO/ANSI LABEL	NO LABEL	UNKNOWN LBL TYPE	AVG MB PER VOL	< MAX CAPACITY
3420	232	215	17	0	82	0	0	71	79	17.5	13(*)
348X	8677	6226	2451	2571	8670	0	0	5	2	80.6	0
TOTAL	8909	6441	2468	2571	8752	0	0	76	81	78.9	13(*)

(*) DATA ON THESE VOLUME(S) EXCEEDED THE MAXIMUM DEVICE CAPACITY **6**
 - THE 'FORCECAP' OPTION CAUSED SOME RECFM=V, AND RECFM=U DATA SETS TO BE RESET TO THE MAXIMUM BYTES PER VOLUME

```
*****
*
* THE FOLLOWING "VOLUME SIZE RANGE TABLE" REPRESENTS THE CURRENT PATTERN OF VOLUME UTILIZATION FOR THE PARTICULAR SET
* OF DATA IN THE SMF SAMPLE SELECTED. IT DOES NOT REPRESENT THE ACTUAL DATA THAT IS CURRENTLY ON THOSE VOLUMES TODAY.*
* SOME OF THE DATA SETS ON THESE CARTRIDGES MAY HAVE ALREADY EXPIRED. EXPIRED VOLSERS ARE INDICATED IN THE VOLUME
* REPORT. THE PURPOSE OF THIS TABLE IS TO SHOW THE CURRENT TAPE UTILIZATION YOU ARE GETTING FOR THE PARTICULAR SET
* OF DATA THAT YOU HAVE FILTERED ON IN THE ASSOCIATED SAMPLE OF SMF INPUT. IT WILL ALSO GIVE YOU A CLEAR INDICATION
* OF HOW MANY TAPE VOLSERS ARE BEING USED BY THAT APPLICATION WITHIN A CERTAIN PERIOD OF TIME.
*
*****
```

Figure 17. GFTASRT2 Summary Report—Tape Volser Processing Phase (The # character in Volume Mount Analyzer reports represents the word "number").

1 The total number of unique volume serials (UNIQUE VOLSERS) is a total volume count taken from the volume perspective (a count of all unique tape volume serials). In this example, the total is 8909.

The GFTASRT2 volume counts are the total gross volume counts. If you run GFTAVMA against one single-file, single-volume data set that was created five times on five different volumes, the GFTASRT2 volume count is five. However, the GFTASRT3 statistical volume count for the data set is one, the current version of the data set. See "Statistical Volumes" on page 31 for more details.

2 Single-file volumes contain a single data set or part of a single data set. See "Single/Multi Volume/File Data Sets" on page 29 for more information.

3 Multifile volumes contain more than one data set. The data sets might be complete, or they might be pieces of data sets that are continued on other volumes. See "Single/Multi Volume/File Data Sets" on page 29 for more information.

4 Returned Scratch volumes are volumes on which all data set copies are obsolete. GFTAVMA has found a more recent copy of every data set, and these volumes are eligible to return to the scratch pool.

To determine the number of scratch volumes, GFTAVMA searches for the most current version of a data set and assumes that its other versions are on volumes returned to scratch. For example, Volume 1 contains data set A on Monday, but Volume 2 contains data set A on Wednesday. Thus, GFTAVMA determines that Volume 1 has been returned to scratch and counted it as a scratch volume.

5 GFTAVMA obtains the average number of megabytes per volume from the information in the SMF records, which were generated at the time the data set was closed. This information does not reflect the effect of IDRC compaction. So, for devices with the IDRC feature, the number of megabytes are non-compacted megabytes. However, any software compaction, such as DFSMSHsm compaction, occurs in the processor and is reflected in the SMF data.

6 **FORCECAP** sets a maximum capacity of data per volume based on cartridge or volume size: 416MB for 3490E cartridges, 208MB for 3490 and 3480 cartridges, and 169MB for 3420 tape volumes. See “The Importance of FORCECAP” on page 141 for more information.

GFTASRT2—Tape Volume Size Range Report

GFTASRT2 also shows your tape library composition. If you specify **CHART** (see “CHART Keyword” on page 134), the raw numbers from this report are put out to the sequential data set defined by the VMACHART DD statement.

Note: You should not try to match these numbers to those from the data set summary report (GFTASRT3) because they simply do not equate. The numbers in GFTASRT2 are from the volume perspective while the numbers in GFTASRT3 are from the data set perspective.

Figure 18 on page 56 shows the tape volume sizes for this run. In this report, keywords were not specified.

GFTASRT2 -- TAPE VOLUME SIZE RANGES									
VOL SIZE	#UNIQUE VOLSERS	% TOT	CUM # VOLSERS	% TOT	# MNTS/ VOLSER	% TOT	CUM # MOUNTS	% TOT	
= 0 MB	230	2.6	230	2.6	871	2.3	871	2.3	
<= 1 MB	2763	31.0	2993	33.6	12320	32.3	13191	34.6	7
<= 5 MB	808	9.1	3801	42.7	3385	8.9	16576	43.5	8
<= 10 MB	597	6.7	4398	49.4	2619	6.9	19195	50.4	9
<= 25 MB	783	8.8	5181	58.2	3300	8.7	22495	59.0	10
<= 40 MB	413	4.6	5594	62.8	1870	4.9	24365	63.9	
<= 50 MB	170	1.9	5764	64.7	711	1.9	25076	65.8	
<= 75 MB	516	5.8	6280	70.5	2215	5.8	27291	71.6	
<=100 MB	458	5.1	6738	75.6	1922	5.0	29213	76.7	
<=125 MB	266	3.0	7004	78.6	1129	3.0	30342	79.6	
<=150 MB	163	1.8	7167	80.4	837	2.2	31179	81.8	
<=175 MB	227	2.5	7394	83.0	935	2.5	32114	84.3	
<=200 MB	130	1.5	7524	84.5	512	1.3	32626	85.6	
<=300 MB	793	8.9	8317	93.4	3070	8.1	35696	93.7	
<=400 MB	348	3.9	8665	97.3	1407	3.7	37103	97.4	
<=500 MB	87	1.0	8752	98.2	306	0.8	37409	98.2	
<=600 MB	83	0.9	8835	99.2	392	1.0	37801	99.2	
<=700 MB	32	0.4	8867	99.5	112	0.3	37913	99.5	
<=800 MB	1	0.0	8868	99.5	2	0.0	37915	99.5	
<=900 MB	12	0.1	8880	99.7	33	0.1	37948	99.6	
<= 1 GB	10	0.1	8890	99.8	66	0.2	38014	99.8	
<= 2 GB	16	0.2	8906	100.0	84	0.2	38098	100.0	
<= 3 GB	1	0.0	8907	100.0	4	0.0	38102	100.0	
<= 4 GB	2	0.0	8909	100.0	6	0.0	38108	100.0	
<= 5 GB	0	0.0	8909	100.0	0	0.0	38108	100.0	
> 5 GB	0	0.0	8909	100.0	0	0.0	38108	100.0	

Figure 18. GFTASRT2 Summary Report—Tape Volume Size Ranges

This report shows you how efficiently you are filling your tape cartridges. Determine the efficiency level of your tape usage, such as 50MB, and use this report to see if you are reaching that level. This example shows that the tape is not efficiently used.

7 This example shows that 2763 unique volume serials contain 1MB of data or less.

8 Forty-two and seven-tenths percent of all of the tapes contain 5MB of data or less. These volumes represent 43.5% of all mounts.

9 This report also shows that 49.4% of the unique volume serials contain 10MB or less. Fifty and four-tenths percent of the mounts at this installation are caused by this inefficient use of tape.

10 Fifty-eight and two-tenths percent of all volumes contain 25MB of data or less. This is only 12% of a native 3480 cartridge capacity or 1% of the capacity that

can be achieved with 3490E and the Enhanced Capacity Cartridge System Tape. In this sample, these volumes represent 59.0% of all mounts.

11 For any given row of the table, CUM # VOLSERS shows the cumulative number of unique volumes that has, at most, that particular amount of data written to them. It is the number of tape volumes used in your library during the period of analysis.

12 For any given row of the table, CUM # MOUNTS shows the number of mounts associated with that number of tape volumes.

Data Set Analysis Phase—GFTASRT3

During the data set analysis phase, GFTAVMA analyzes the data from the data set perspective and generates the GFTASRT3 summary report. If data set 'X' was written ten times during the sample, GFTAVMA considers the last time it was written to be the current version of data set 'X'. GFTAVMA attributes all the mounts caused by all versions of data set 'X' to data set 'X', but GFTASRT3 shows only the portion of statistical volumes attributed to the latest copy of the data set. See "Statistical Volumes" on page 31 for an explanation of unique volsers and statistical volumes.

The GFTAVMA keywords **DATE** or **TIME** change what GFTAVMA shows as the current version of data set 'X'.

GFTASRT3—Tape Data Set Statistics Report

GFTASRT3 is the part of the GFTAVMA summary report that shows the current version of a data set, its current size, the job and program that created it, and all mounts attributed to all versions of that data set. It shows only the current volume the data set resides on.

If you specify the **GDG(GROUP)** keyword, the numbers in the Tape Data Set Statistics Report reflect non-GDG data sets and GDG data sets that are grouped together and counted as one generation data group.

Figure 19 on page 58 is the report of tape data set statistics. In this report, keywords were not specified.

GFTASRT3 Report

GFTASRT3 -- TAPE DATA SET PROCESSING PHASE -- 05/21/2004 -- 03:45:10 P.M.

```
*****
*
*           GFTASRT3 TAPE DATA SET STATISTICS           *
*
*****
```

```
38702   GDS(S) - ( 93.9%) --    2049 GDG GROUP(S)    200 MAX GDS/GDG  1
2536   NON-GDS(S) - (  6.1%)
```

```
-----
41238(+) TOTAL - (100.0%)    40612 IDRC DS ( 98.5%)
```

(+) 2 OF THESE DATA SET(S) HAD A CALCULATED SIZE WHICH EXCEEDED THE MAXIMUM VOLUME CAPACITY
-- THE SIZE WAS OVERRIDDEN DUE TO THE 'FORCECAP' KEYWORD

```
31.7 MB (AVG PER DATA SET) 2 38584.8 MOUNTS (STATISTICAL) 5
1309.1 GB (CUM TOTAL ALL DS) 3 6338.1 VOLUMES (STATISTICAL) 6
31 DAYS (MAX REF AGE) 4 0.3 DAYS (AVG MAX REF AGE) 7 8
```

SINGLE-VOLUME DATA SET(S)									
DATA SET FILE TYPE	# OF DSNS	% OF DSNS	# OF GB	% OF GB	# OF STAT MOUNTS	% STAT MOUNTS	# OF STAT VOLUMES	% STAT VOLUMES	
SING-FILE	9474	23.0%	276.1	21.1%	21629.2	56.1%	3626.0	57.2%	9
MULT-FILE	29096	70.6%	328.6	25.1%	10922.3	28.3%	1786.0	28.2%	

MULTIVOLUME DATA SET(S)									
DATA SET FILE TYPE	# OF DSNS	% OF DSNS	# OF GB	% OF GB	# OF STAT MOUNTS	% STAT MOUNTS	# OF STAT VOLUMES	% STAT VOLUMES	
SING-FILE	469	1.1%	352.1	26.9%	4603.1	11.9%	619.0	9.8%	
MULT-FILE	2199	5.3%	352.3	26.9%	1430.1	3.7%	307.1	4.8%	
TOTAL DSN	41238	100.0%	1309.1	100.0%	38584.8	100.0%	6338.1	100.0%	

```
*****
*
* THE FOLLOWING "DATA SET SIZE RANGE TABLE" REPRESENTS THE CURRENT PATTERN
* OF DATA SET SIZES FOR THE PARTICULAR SET OF DATA IN THE SMF SAMPLE
* SELECTED. IT DOES NOT IMPLY THAT ALL OF THESE DATA SETS STILL EXIST.
* SOME OF THEM MAY HAVE ALREADY EXPIRED. THE TABLE DOES SHOW THE SIZE OF
* EACH DATA SET THE LAST TIME THAT DATA SET HAD BEEN WRITTEN (IF KNOWN).
* THE PURPOSE OF THIS TABLE IS TO SHOW THE CURRENT SIZE RANGES OF DATA SETS
* THAT ARE BEING CREATED IN YOUR INSTALLATION ON THE PARTICULAR SET OF DATA
* YOU HAVE FILTERED ON. THESE SIZE RANGES SHOULD CLOSELY MATCH THE SIZE
* RANGES OF THE DATA SETS THAT ACTUALLY DO EXIST ON YOUR SYSTEM. THE MOST
* RECENT SIZE OF THE DATA SET IS ALWAYS USED.
*
*****
```

Figure 19. GFTASRT3 Summary Report—Tape Data Set Processing Phase

1 The GDG statistics show the number of GDSs per GDG during the sample period.

- 2** The size of data sets on tape averaged 31.7MB for this run.
- 3** The total amount of data represented by the 41,238 data sets in this run is 1309.1GB.
- 4** The age of the oldest data set in the sample, where age is the number of days between the date of last reference and the end of the sample, is 31 days. If this value is less than the number of days in the sample, all data sets from the early part of the sample have been deleted or referenced during the sample period.
- 5** For an explanation of statistical mounts, see “Statistical Mounts” on page 30.
- 6** For an explanation of statistical volumes, see “Statistical Volumes” on page 31.
- 7** The average of all the maximum numbers of days since last reference for all data sets in the sample. If this number is significantly less than 1, such as 0.1, it indicates that the majority of the data sets in this sample are referenced as backup and can be migrated off the tape mount management DASD buffer hourly.
- 8** This line does not print if you specify **GDG(GROUP)**.
- 9** Note the high percentage of single file data sets (23.0%) and the percentage of mounts they represent (56.1%). These data sets are probably data sets that were created for backup purposes during batch processing. They are all necessary, but they contribute to inefficient use of tape. If you manage only these data sets with tape mount management, you can reduce the majority (56.1%) of your tape mounts.

GFTASRT3—Tape Data Set Size Range Report

GFTASRT3 shows the most current version of each particular data set found during the sample. If ten copies of a data set exist on tape, this part of the report only includes the latest version.

This report shows the number of data sets in each size range and the mount activity that is attributed to the data sets in that range. If you specify **CHART** (see “CHART Keyword” on page 134), the raw numbers from this report are put out to the sequential data set defined by the VMACHART DD statement.

If you specify the **GDG(GROUP)** keyword, the numbers in the Tape Data Set Size Range Report reflect non-GDG data sets and GDG data sets grouped together and counted as one generation data group.

Note: You should not try to match these numbers to those from the volume summary report (GFTASRT2) because they simply do not equate. The numbers in GFTASRT3 are from the data set perspective while the numbers in GFTASRT2 are from the volume perspective.

Figure 20 on page 60 is the report that summarizes data set sizes for this run. In this report, keywords were not specified.

GFTASRT3 Report

10a		10b											
GFTASRT3 -- TAPE DATA SET SIZE RANGES													
DSN SIZE	# OF DSNS	% TOT DSNS	CUM # DSNS	% TOT DSNS	# OF STAT MOUNTS	% TOT MNTS	CUM STAT MOUNTS	% TOT MNTS	# OF STAT VOLUMES	% TOT VOLS	CUM STAT VOLUMES	% TOT VOLS	
= 0 MB	2568	6.2	2568	6.2	1277.8	3.3	1277.8	3.3	192.6	3.0	192.6	3.0	11
<= 1 MB	21368	51.8	23936	58.0	15173.7	39.3	16451.5	42.6	2480.8	39.1	2673.4	42.2	12
<= 5 MB	4870	11.8	28806	69.9	3532.6	9.2	19984.1	51.8	667.8	10.5	3341.2	52.7	
<= 10 MB	2203	5.3	31009	75.2	2125.0	5.5	22109.1	57.3	412.5	6.5	3753.6	59.2	12
<= 25 MB	2852	6.9	33861	82.1	3379.6	8.8	25488.7	66.1	524.6	8.3	4278.2	67.5	
<= 40 MB	1888	4.6	35749	86.7	1877.3	4.9	27366.0	70.9	345.3	5.4	4623.5	72.9	
<= 50 MB	691	1.7	36440	88.4	670.6	1.7	28036.5	72.7	122.2	1.9	4745.7	74.9	
<= 75 MB	1382	3.4	37822	91.7	1828.4	4.7	29865.0	77.4	274.8	4.3	5020.5	79.2	
<=100 MB	710	1.7	38532	93.4	991.4	2.6	30856.3	80.0	143.1	2.3	5163.6	81.5	13
<=125 MB	502	1.2	39034	94.7	708.0	1.8	31564.3	81.8	94.0	1.5	5257.6	83.0	
<=150 MB	141	0.3	39175	95.0	195.7	0.5	31760.0	82.3	37.4	0.6	5295.0	83.5	
<=175 MB	245	0.6	39420	95.6	350.5	0.9	32110.6	83.2	68.6	1.1	5363.6	84.6	
<=200 MB	335	0.8	39755	96.4	357.0	0.9	32467.6	84.1	71.9	1.1	5435.5	85.8	
<=300 MB	586	1.4	40341	97.8	768.4	2.0	33236.0	86.1	146.4	2.3	5581.9	88.1	
<=400 MB	257	0.6	40598	98.4	587.0	1.5	33823.0	87.7	127.2	2.0	5709.1	90.1	
<=500 MB	49	0.1	40647	98.6	203.3	0.5	34026.4	88.2	20.3	0.3	5729.3	90.4	
<=600 MB	174	0.4	40821	99.0	547.4	1.4	34573.7	89.6	120.8	1.9	5850.1	92.3	
<=700 MB	21	0.1	40842	99.0	466.6	1.2	35040.4	90.8	50.0	0.8	5900.1	93.1	
<=800 MB	30	0.1	40872	99.1	512.0	1.3	35552.4	92.1	78.0	1.2	5978.1	94.3	
<=900 MB	57	0.1	40929	99.3	760.0	2.0	36312.4	94.1	100.0	1.6	6078.1	95.9	
<= 1 GB	84	0.2	41013	99.5	829.3	2.1	37141.6	96.3	114.6	1.8	6192.7	97.7	
<= 2 GB	216	0.5	41229	100.0	1381.1	3.6	38522.8	99.8	137.4	2.2	6330.1	99.9	
<= 3 GB	6	0.0	41235	100.0	50.0	0.1	38572.8	100.0	4.0	0.1	6334.1	99.9	
<= 4 GB	1	0.0	41236	100.0	4.0	0.0	38576.8	100.0	2.0	0.0	6336.1	100.0	
<= 5 GB	1	0.0	41237	100.0	4.0	0.0	38580.8	100.0	2.0	0.0	6338.1	100.0	
> 5 GB	1	0.0	41238	100.0	4.0	0.0	38584.8	100.0	0.0	0.0	6338.1	100.0	

GFTAVMA PROGRAM COMPLETED AT 03:45:37 P.M. ON 05/21/2004 -- RETURN CODE = 0

Figure 20. GFTASRT3 Summary Report—Tape Data Set Analysis

This report shows the sizes of your data sets and the number of mounts these data sets caused. If small data sets are causing most of your mounts, as they are in this example, you are using your tape media inefficiently.

10 Compare the number of data sets in each of the size ranges. Notice that the majority of the data sets are in the smaller size ranges while fewer data sets are in the large size ranges. The data set size (DSN SIZE) and number of mounts (CUM STAT MOUNTS and % TOT MNTS) show that the smaller data sets caused most of the mounts in this example.

11 Two-thousand-five-hundred-sixty-eight data sets are equal to 0MB and account for 1277.8 mounts. These data sets were opened, but had no read or write activity. These data sets should not have caused tape mounts because there was no transfer of data. These mounts could have been avoided by directing these data sets to the tape mount management DASD buffer.

12 This sample is typical in that the majority of data sets on tape are small, yet they account for most of the mounts. For example, 75.2% of the data sets analyzed are less than or equal to 10MB, but they account for 57.3% of all tape mounts in this run.

13 Small data sets, less than or equal to 100MB, are causing 80% of the mounts in this example, which indicates that tape mount management could significantly reduce the number of mounts for this installation.

Chapter 5. Using GFTAVMA Keywords to Filter Your Input

This chapter describes the GFTAVMA keywords you use to filter the input to the volume mount analyzer. As introduced in “Filtering Your Input with GFTAVMA Keywords” on page 44, these keywords affect the input phase of processing by selecting the SMF records the volume mount analyzer analyzes.

All GFTAVMA keywords are optional. The VMACNTL DD statement points to a sequential data set with 80-byte records that contain the GFTAVMA keywords. Normally, this is a DD * data set. See the GFTAVMA JCL member in “Submitting the GFTAVMA JCL” on page 48.

In this chapter, the keywords are arranged in alphabetical order. Sub-keywords, if any, are described with their corresponding keywords.

Keyword Rules

When using keywords to filter your input, you should:

- Start or end keywords in any position between 1 and 71. Positions 72 through 80 are ignored
- Separate keywords with either blanks or commas
- Separate sub-keywords with commas
- Separate items in a list with commas
- Put keywords in any order
- Use the correct number of parentheses

If a keyword spans more than one record, you must indicate continuation by using continuation characters in the last text character of the keyword. Blanks can follow a continuation character in a record. The valid continuation characters for GFTAVMA are a comma and the left and right parenthesis.

Comments are surrounded by the delimiters /* and */. A comment start delimiter, /*, cannot be in column one. Examples:

```
TIME(0,12) /* This is a comment */
TIME(0,   /* The last character, the comma, indicates continuation */
  12)    /* part two of the continued keyword */
```

Summary of Keywords That Filter Your Input

Table 5 summarizes the filtering keywords. These keywords affect which SMF records are used as input to GFTAVMA.

Table 5. Keywords that Filter Your Input

Keyword	Filter Type ¹	Page Number
ACCOUNT	Primary Include/Exclude	64
DATASET	Primary Include/Exclude	65
DATE	Limiting	66
EXPDT	Primary Include/Exclude	67
FILE	Limiting	68
JOBNAME	Primary Include/Exclude	69

Table 5. Keywords that Filter Your Input (continued)

Keyword	Filter Type ¹	Page Number
MAXSIZE	Limiting	70
MINSIZE	Limiting	70
MOUNT	Limiting	72
PROGRAM	Primary Include/Exclude	73
SYSTEMID	Secondary Include/Exclude	74
TIME	Limiting	74
UNIT	Secondary Include/Exclude	75
UNITADDR	Secondary Include/Exclude	76
USAGE	Limiting	78
VOLUME	Secondary Include/Exclude	79

Note:

1. See “Filtering Your Input with GFTAVMA Keywords” on page 44 for complete details on filtering and filtering keywords.

See Appendix A, “A Quick Reference to Volume Mount Analyzer Keywords,” on page 157 for a complete list of the keyword abbreviations and defaults. See Appendix B, “Volume Mount Analyzer Keyword Syntax,” on page 159 for the syntax of all keywords.

ACCOUNT Keyword

Use **ACCOUNT** to include or exclude data sets created by jobs that have specified accounting information. If a data set is not created during the time period of the SMF sample, the **ACCOUNT** information is from the first reference to the data set. See “Primary Filters” on page 45 for more information.

The accounting information that GFTAVMA uses is in the first 8 bytes of the accounting field as specified on the JOB statement.

Abbreviation: ACCT

Syntax

Figure 21 shows the syntax for the **ACCOUNT** keyword.

[**ACCOUNT (INCLUDE** (*account number include list*),
EXCLUDE (*account number exclude list*))]

Figure 21. ACCOUNT Keyword Syntax

INCLUDE (*account number include list*)

Specifies a list of account numbers to include in the analysis.

EXCLUDE (*account number exclude list*)

Specifies a list of account numbers to exclude from the analysis.

See “Include/Exclude Filter Syntax” on page 45 for the rules and syntax for filters.

Default

If you do not specify the **ACCOUNT** keyword, GFTAVMA analyzes all data sets.

Relationship to Other Keywords

There can be more than one accounting field in the JOB accounting information. The first accounting field is the default. To override this default, use the **ACCTFLD** keyword when you run GFTAXTR to specify another accounting field. See “ACCTFLD Keyword” on page 15 for more information.

ACCOUNT Examples

If you specify:

```
ACCOUNT (INCLUDE (%7*),
        EXCLUDE (1*))
```

GFTAVMA includes all data sets that were created by jobs that have account numbers with a 7 in the second position but excludes any data sets that were created by jobs that have account numbers that start with 17.

Using ACCOUNT as an ACS Filter

ACCOUNT is a filter you can use for automatic class selection (ACS) routines. GFTAVMA analyzes only eight bytes of the account field, while the **ACCOUNT** parameter for ACS is the value of the entire JOB account field.

DATASET Keyword

Use **DATASET** to include or exclude data sets from the GFTAVMA analysis based on the data set names in the input SMF data records. See “Primary Filters” on page 45 for more information.

Abbreviation: DSN

Syntax

Figure 22 shows the syntax for the **DATASET** keyword.

```
[DATASET (INCLUDE (data set filter list),
         EXCLUDE (data set filter list))]
```

Figure 22. DATASET Keyword Syntax

INCLUDE (*data set filter list*)

Specifies a list of data sets to include in the analysis.

EXCLUDE (*data set filter list*)

Specifies a list of data sets to exclude from the analysis.

See “Include/Exclude Filter Syntax” on page 45 for the rules and syntax for filters.

Default

If you do not specify the **DATASET** keyword, GFTAVMA will analyze all the data sets.

DATASET Examples

If you specify:

```
DATASET (INCLUDE (IMS.** , DB2.**),
        EXCLUDE (IMS.IMAGE.COPY.**))
```

GFTAVMA includes all the IMS/ESA[®] and DB2[®] database data sets except those known from their data set names as image copies for IMS.

If you specify:

```
DATASET (EXCLUDE (DFDSS.WEEKLY.DUMP.** , DFDSS.DR.**))
```

GFTAVMA excludes all the input SMF data associated with *DFSMSDss*[™] weekly dumps and DFSMSDss error-recovery dumps based on the data set naming rules for the output data sets.

Using DATASET as an ACS Filter

DATASET is a filter you can use for the ACS routines. However, data set names are not the most useful ACS filter because listing the data set names makes your ACS filter list long and difficult to maintain.

DATE Keyword

Use **DATE** to select SMF records created on the specified dates or in the specified date range. See “Limiting Filters” on page 44 for more information.

Abbreviation: None

Syntax

Figure 23 shows the syntax for the **DATE** keyword.

[DATE (*fromdate*[,*todate*])]

Figure 23. DATE Keyword Syntax

fromdate

Specifies the start date for the records to include in your analysis.

todate

Specifies the end date for the records to include in your analysis.

If you specify only the *fromdate*, GFTAVMA selects only the records created on that date.

You can use any of the following formats for *fromdate* and *todate*:

- | | |
|-------------------|--|
| mm/dd/yy | The month/day/year specified with 2 digit fields, for example:
03/15/04 or 3/15/04 |
| mm/dd/yyyy | The month/day/year where the month and day are 2 digit fields and the year is a 4 digit field, for example: 03/15/2004 |
| yyyyddd | The year and Julian date using 4 digits for the year, for example:
2005074 |
| yyddd | The year and Julian date using 2 digits for the year, for example:
05074 |

Default

If you do not specify the **DATE** keyword, GFTAVMA analyzes the SMF records for every day in the collected sample.

DATE Examples

If you specify:

DATE (05/30/05)

GFTAVMA includes only SMF records that were created on May 30, 2005.

If you specify:

DATE (05/27/00,06/28/00)

GFTAVMA includes all SMF records created from May 27, 2000, through June 28, 2000.

EXPDT Keyword

Use **EXPDT** to include or exclude data sets from the analysis based on the expiration date of the creating job's data set. If the data set was not created during the time period of the SMF sample, the **EXPDT** information is from the first reference to the data set. See "Primary Filters" on page 45 for more information.

If a data set does not have a specified expiration date, it is not included when you specify **EXPDT (INCLUDE)**. Only the data sets whose expiration dates match the **EXPDT** include list are included in the analysis. You cannot explicitly include or exclude data sets with blank expiration dates.

The usefulness of this function varies depending on the way that your installation uses expiration dates. For example, some sites use a specific value, such as 04000, in the expiration date field to indicate a data set that is to go off site after its creation. This use provides a simple approach to identify and inform the Storage Management Subsystem (SMS) those data sets that must be allocated directly to tape.

Abbreviation: EXP

Syntax

Figure 24 shows the syntax for the **EXPDT** keyword.

```
[EXPDT (INCLUDE (nnnnn),
        EXCLUDE (yyyyyy))]
```

Figure 24. EXPDT Keyword Syntax

INCLUDE (nnnnn)

Specifies a list of expiration dates to include in the analysis.

EXCLUDE (yyyyyy)

Specifies a list of expiration dates to exclude from the analysis.

See "Include/Exclude Filter Syntax" on page 45 for the rules and syntax for filters.

EXPDT

Use the following date formats for **EXPDT**:

yyyyddd	The year and Julian date using 4 digits for the year, for example: 2005074
yyddd	The year and Julian date using 2 digits for the year, for example: 05074

Default

If you do not specify the **EXPDT** keyword, GFTAVMA analyzes all data sets.

EXPDT Examples

If you specify:

```
EXPDT (INCLUDE (05*), EXCLUDE (05005))
```

GFTAVMA includes data sets with an expiration date matching the filter *05** (the year 2005) except those that have an expiration date of *05005* (January 5, 2005).

Using EXPDT as an ACS Filter

EXPDT is a filter you can use for the automatic class selection (ACS) routines.

FILE Keyword

Use **FILE** to select:

- Data sets that reside alone on one or more volumes
- Data sets that reside with other data sets on one or more volumes
- All data sets

See “Limiting Filters” on page 44 for more information.

Abbreviation: None

Syntax

Figure 25 shows the syntax for the **FILE** keyword.

[FILE (SINGLE MULTIPLE <u>BOTH</u>)]

Figure 25. FILE Keyword Syntax

SINGLE

Includes those volumes that contain a single data set. All volumes that have multiple data sets are excluded.

MULTIPLE

Includes those volumes that contain multiple data sets. All volumes that have only a single data set are excluded.

BOTH

Includes all volumes in the analysis.

Single-file and multifile data sets can be either single-volume or multivolume. See “Single/Multi Volume/File Data Sets” on page 29 for more information about tape data sets.

Default

The default for FILE is BOTH.

Relationship to Other Keywords

If you want to examine the effects of IDRC compaction, 3490E, or Enhanced Capacity Cartridge System Tape on your large, multivolume data sets, such as volume or database dumps, use **FILE(SINGLE)** and **MOUNT(MULTIPLE)**. This keyword combination causes GFTAVMA to filter in those large, multivolume data sets.

FILE Examples

If you specify:

FILE (BOTH)

GFTAVMA will include all volumes regardless of the number of data sets that reside on them.

If you specify:

FILE (SINGLE)

GFTAVMA will include only those volumes that contain a single data set.

Note: If a volume is used more than once, the last data written to the volume is considered to be what is on the volume.

JOBNAME Keyword

Use **JOBNAME** to include or exclude data sets created by specific job names. All data sets associated with the execution of a job are included in or excluded from the analysis. If you code **INCLUDE**, then all references to those data sets are included, including references from other jobs unless they are excluded by another filter. If a data set was not created during the time period of the SMF sample, **JOBNAME** is from the first reference to the data set. See “Primary Filters” on page 45 for more information.

Abbreviation: JOB

Syntax

Figure 26 shows the syntax for the **JOBNAME** keyword.

<p>[JOBNAME (INCLUDE (job name filter list), EXCLUDE (job name filter list))]</p>
--

Figure 26. **JOBNAME** Keyword Syntax

INCLUDE (job name filter list)

Specifies a job name filter list to include in the analysis.

EXCLUDE (job name filter list)

Specifies a job name filter list to exclude from the analysis.

See “Include/Exclude Filter Syntax” on page 45 for the rules and syntax for these keywords.

Default

If you do not specify the **JOBNAME** keyword, GFTAVMA analyzes data sets associated with all job names.

JOBNAME Example

If you specify:

```
JOBNAME (EXCLUDE (IPCSOMP, SMF*, DLOG))
```

GFTAVMA will exclude all data sets for the jobs named IPCSDMP, SMF*, and DLOG. GFTAVMA includes data sets from all other jobs in the analysis.

Using JOBNAME as an ACS Filter

JOBNAME is a filter you can use for the ACS routines.

MAXSIZE and MINSIZE Keywords

Use **MAXSIZE** to analyze data sets that are less than or equal to **MAXSIZE**. Similarly, use **MINSIZE** to analyze data sets that are greater than or equal to **MINSIZE**. See “Limiting Filters” on page 44 for more information.

Abbreviations: **MAX** and **MIN**

Syntax

Figure 27 shows the syntax for the **MAXSIZE** and **MINSIZE** keywords.

[MAXSIZE (nnnn) MINSIZE (nnnn)]

Figure 27. **MAXSIZE** and **MINSIZE** Keyword Syntax

MAXSIZE (nnnn)

Includes all data sets that are less than or equal to *nnnn* megabytes in the analysis.

MINSIZE (nnnn)

Includes all data sets that are greater than or equal to *nnnn* megabytes in the analysis.

You can specify either **MAXSIZE** or **MINSIZE** or both for a run. The ranges for **MAXSIZE** and **MINSIZE** are zero to infinity.

Default

The default for **MAXSIZE** is infinity, and the default for **MINSIZE** is zero. If you do not specify the **MAXSIZE** and **MINSIZE** keywords, GFTAVMA analyzes all data sets.

Caution

Do not use **MAXSIZE** and **MINSIZE** to set artificial bounds on the data to be managed by the Storage Management Subsystem (SMS). In other words, do not use GFTAVMA to identify all data sets that are less than 10MB. In this way, you can move these data sets to DASD. This approach is extremely difficult to implement and often results in thousands of ACS filters.

By restricting the size of data sets analyzed in this manner, you do not get meaningful patterns to develop filters for ACS routines. Therefore, in order to capture that set of data, you end up with large numbers of fully-qualified data set names in the ACS routines, a situation that quickly becomes cumbersome and unmanageable.

Note: **MAXSIZE** and **MINSIZE** ignore any data sets that do not fall into the selected size range. If one generation of a GDG happens to fall outside the selected size range, it is not considered in the analysis even though all other generations in that GDG are included.

Do not be too dependent on **MAXSIZE** and **MINSIZE** because they do not offer an easy method of tape mount management implementation. Most tape data sets do not have size coded in the JCL. By the time the actual size of the data set is determined, it is too late to influence its fate with regard to being intercepted by DFSMS or not.

The size of a data set can certainly be a common denominator if it is looked at as an average or range rather than an absolute number. For instance, many large data sets probably have some common characteristics that emerge as patterns when you analyze them. Certain programs or jobs tend to generate many of the large data sets. However, there is no number that is a hard and fast breakpoint as to what determines a large data set.

Use **MAXSIZE** and **MINSIZE** as a first cut, but in the final analysis, you need to develop filters based on attributes that can be used by ACS to intercept the data sets. Otherwise, you do not have a practical basis for tape mount management implementation, and your study has been an interesting paper exercise at best.

Relationship to Other Keywords

The default value for the **LARGE** keyword is 600MB, but users have successfully intercepted data sets as large as 1.8GB. Using **GFTAVMA**, you can determine useful filters for **LARGE** using **MINSIZE**. If you specified **MINSIZE (600)** and **REPORT (DSN)**, you may have 30 data sets for ABC.PROD that are over 600MB. By reducing the **MINSIZE** value to 400MB, the report might show 2200 data sets for ABC.PROD.

MAXSIZE and MINSIZE Examples

If you specify:

```
MINSIZE (600)
```

GFTAVMA will include only large data sets that are candidates to go directly to tape with hardware compaction. You set 600MB as the minimum size because that is the approximate capacity of a cartridge using IDRC compaction or half the capacity of a 3490E Enhanced Capacity Cartridge System Tape. With 3590, you will want to specify a larger value.

If you specify:

```
MAXSIZE (600)
```

GFTAVMA will include data sets that are smaller than 600MB. You can use 600MB as the estimated maximum size to analyze only the data sets that you would consider candidates to intercept and manage with DFSMS.

MOUNT Keyword

Use **MOUNT** to select:

- Data sets that are contained on one volume
- Only those data sets that span multiple volumes
- All data sets regardless of the number of residence volumes

See “Limiting Filters” on page 44 for more information.

Abbreviation: MOU

Syntax

Figure 28 shows the syntax for the **MOUNT** keyword.

`[MOUNT (SINGLE | MULTIPLE | BOTH)]`

Figure 28. MOUNT Keyword Syntax

SINGLE

Includes only those data sets that are contained on a single volume. All multiple volume data sets are excluded.

MULTIPLE

Includes only those data sets that span multiple volumes. All single volume data sets are excluded.

BOTH

Includes all data sets in the analysis.

Each of the above types of data sets can be either single-file or multifile. See “Single/Multi Volume/File Data Sets” on page 29 for more information about tape data sets.

Default

The default for **MOUNT** is **BOTH**.

Relationship to Other Keywords

If you want to examine the effects of IDRC on your large, multivolume data sets, such as volume or database dumps, use **FILE(SINGLE)** and **MOUNT(MULTIPLE)**. This keyword combination causes GFTAVMA to filter in those large, multivolume data sets.

MOUNT Examples

If you specify:

```
MOUNT (BOTH)
```

GFTAVMA will include all data sets regardless of their volume residence requirements.

If you specify:

```
MOUNT (MULTIPLE)
```

GFTAVMA will include only those data sets that reside on more than one volume.

PROGRAM Keyword

Use **PROGRAM** to include or exclude data sets created by specific program names. All data sets created by the named programs are included in or excluded from the data analyzed. If a data set was not created during the time period in the SMF sample, the **PROGRAM** information is from the first reference to the data set. See “Primary Filters” on page 45 for more information.

Abbreviation: PGM

Syntax

Figure 29 shows the syntax for the **PROGRAM** keyword.

```

PROGRAM (INCLUDE (program filter list),
EXCLUDE (program filter list))]
```

Figure 29. **PROGRAM** Keyword Syntax

INCLUDE (program filter list)

Specifies a program filter list to include in the analysis.

EXCLUDE (program filter list)

Specifies a program filter list to exclude from the analysis.

See “Include/Exclude Filter Syntax” on page 45 for the rules and syntax for these keywords.

Default

If you do not specify the **PROGRAM** keyword, GFTAVMA analyzes all data sets associated with all programs.

PROGRAM Examples

If you specify:

```
PROGRAM (INCLUDE (IDCAMS, IEB*, IEH*),
EXCLUDE (ADRSSU))
```

GFTAVMA will include all the data being handled by IBM utilities, except the data handled by DFSMSdss (ADRSSU).

If you specify:

```
PROGRAM (EXCLUDE (IKJEFT01))
```

GFTAVMA will exclude data handled by the TSO Terminal Monitor Program (IKJEFT01).

Using PROGRAM as an ACS Filter

PROGRAM is a filter you can use for the automatic class selection (ACS) routines.

SYSTEMID Keyword

Use **SYSTEMID** to include or exclude particular systems in the analysis. Each system that generates SMF data is identified by a system ID. The **SYSTEMID** keyword provides the ability to focus on one or more specific systems in the analysis. See “Secondary Filters” on page 45 for more information.

Abbreviation: SYID

Syntax

Figure 30 shows the syntax for the **SYSTEMID** keyword.

<pre>[SYSTEMID (INCLUDE (system-id1), EXCLUDE (system-id2))]</pre>
--

Figure 30. **SYSTEMID** Keyword Syntax

INCLUDE (*system-id1*)

Specifies a list of system IDs to include in the analysis.

EXCLUDE (*system-id2*)

Specifies a list of system IDs to exclude from the analysis.

See “Include/Exclude Filter Syntax” on page 45 for the rules and syntax for filters.

Default

If you do not specify the **SYSTEMID** keyword, GFTAVMA will analyze data sets from all systems.

SYSTEMID Example

If you specify:

```
SYSTEMID (EXCLUDE (TEST))
```

GFTAVMA will exclude TEST, which is a logical test partition (LPAR). You would exclude LPARs because test cases created the LPAR data sets, which generated mounts. Since these data sets were created by test cases, you would not intercept and manage them with DFSMS.

TIME Keyword

Use **TIME** to select records created at specified times or within a specified time range. Using **TIME** can be useful for examining a daily processing window or a particular shift. See “Limiting Filters” on page 44 for more information.

If you specify a time range using **TIME** (*fromtime, totime*), only records within that time range are extracted for each day in the sample.

Abbreviation: None

Syntax

Figure 31 on page 75 shows the syntax for the **TIME** keyword.

[TIME (<i>fromtime</i> [, <i>totime</i>])]
--

Figure 31. TIME Keyword Syntax

Both *fromtime* and *totime* have this syntax:

hh:mm:ss

Specifies the hour, minute, and second of the timestamp.

If you only specify the *fromtime*, the default *totime* is the end of the day, 23:59:59.

The valid range of values is between 00:00:00 for midnight, which is the start of the day, and 23:59:59, which is the end of the day.

Times can be specified as hh:mm:ss, hh:mm, or just hh. The time values specified must be in 24 hour clock values.

Default

If you do not specify the **TIME** keyword, GFTAVMA will analyze the SMF records for all times in the sample collected.

TIME Examples

If you specify:

TIME (3)

GFTAVMA includes only records created between 03:00:00 and 23:59:59 on each day of the sample period.

If you specify:

TIME (8,12)

GFTAVMA will include only records created between 08:00:00 and 12:00:00 on each day of the sample period.

UNIT Keyword

Use **UNIT** to select the unit types of tape volumes to be included in the analysis. Specify the unit type to differentiate between hardware devices with varying recording densities and media capacities. See “Secondary Filters” on page 45 for more information.

Since only the actual device type is in the SMF record, GFTAVMA only recognizes actual device types. It does not recognize esoteric unit types.

Abbreviation: None

Syntax

Figure 32 on page 76 shows the syntax for the **UNIT** keyword.

UNIT

```
[UNIT (INCLUDE (3420,3480,348X,3490E,3590),  
        EXCLUDE (3420,3480,348X,3490E,3590))]
```

Figure 32. UNIT Keyword Syntax

INCLUDE (3420,3480,348X,3490E,3590)

Specifies a list of unit types to include in the analysis.

EXCLUDE (3420,3480,348X,3490E,3590)

Specifies a list of unit types to exclude from the analysis.

You can specify one or more unit types. The valid unit types are:

- 3420** 3420 type reel tape
- 3480** 3480 cartridge tape without the IDRC feature
- 348X** 3480 cartridge tape with the IDRC feature or 3490
- 3490E** 3490 cartridge tape written with enhanced recording format and 36-track processing
- 3590** 3590 cartridge tape written with 128-track processing and 36-track processing

The **UNIT** keyword allows the use of global data set name characters:

```
UNIT (INC(348%))
```

See “Include/Exclude Filter Syntax” on page 45 for the rules and syntax for filters.

Default

If you do not specify the **UNIT** keyword, GFTAVMA will analyze data sets on all units.

Relationship to Other Keywords

Because the 3480 with IDRC and the 3490 create the same capacity and format of cartridge, the SMF records do not differentiate between these with separate unit types. To separate these in an GFTAVMA analysis, use the **UNITADDR** keyword.

UNIT Examples

If you specify:

```
UNIT (EXC(3420))
```

GFTAVMA will include data sets created on any device type except 3420s.

UNITADDR Keyword

Use **UNITADDR** to include a specific range of device numbers in the analysis or to exclude a range of devices that are of a different device type. See “Secondary Filters” on page 45 for more information. In the book and in the reports device numbers are called “unit addresses”.

UNITADDR is useful if you use certain address ranges specifically for some particular application or function. For example, you might have several tape units in a locked room that are used only for tapes containing highly confidential data. Therefore, you may want to exclude that range of addresses from your analysis.

Abbreviation: UADR

Syntax

Figure 33 shows the syntax for the **UNITADDR** keyword.

```
[UNITADDR (INCLUDE (xxxx),
           EXCLUDE (yyyy))]
```

Figure 33. **UNITADDR** Keyword Syntax

INCLUDE (xxxx)

Specifies a list of unit addresses to include in the analysis.

EXCLUDE (yyyy)

Specifies a list of unit addresses to exclude from the analysis.

Unit addresses are four characters in length (for example, 0580).

See “Include/Exclude Filter Syntax” on page 45 for the rules and syntax for filters.

Default

If you do not specify the **UNITADDR** keyword, GFTAVMA will analyze data sets from all unit addresses.

Relationship to Other Keywords

When filtering with **VOLUME** and **UNITADDR**, ensure that the filtering combination works. For example, if data set AA is a multivolume data set that resides on volume 1, volume 2, and volume 3, it is possible that when you use **VOLUME** and **UNITADDR** filtering, you can filter volume 2 out of the analysis. This would affect the volume mount analyzer report results for data set AA.

To avoid this situation, only use these keywords to capture a totally self-contained set of volumes, such as DFSMSHsm-owned volumes. These self-contained sets of volumes have a naming convention that identifies them to be used by one particular application only.

Caution

Including or excluding specific unit addresses can sometimes produce inaccurate results. For example, if data set X.Y.Z is multivolume and spans volumes 1, 2, and 3, and if volume 2 gets excluded due to filtering, the data associated with data set X.Y.Z becomes meaningless.

UNITADDR Examples

If you specify:

```
UADR (INCLUDE (07*), EXCLUDE (07C%,07D%))
```

GFTAVMA will include all unit addresses in the range of 0700 - 07FF, except for 07C0-07CF and 07D0-07DF.

If you specify:

```
UNITADDR (INCLUDE (07A*,07B*))
```

GFTAVMA will include all unit addresses in the range of 07A0 to 07BF. You can use an asterisk to specify a range of unit addresses.

USAGE Keyword

Use **USAGE** to select data sets for analysis based on their patterns of usage or recognizable naming conventions. **USAGE** simplifies the categorization of tape data sets and limits the number of filters that you need to code in individual GFTAVMA runs. See “Limiting Filters” on page 44 for more information.

Note that the data set usage attributes are not the same as the tape mount management data set categories. See “Data Set Usage Types” on page 36 for more information.

Abbreviation: USE

Syntax

Figure 34 shows the syntax for the **USAGE** keyword.

```
[USAGE (ALL | ACTIVE, BACKUP, BCOPY, TEMP,
        SINGLE | HSM | NONHSM)]
```

Figure 34. USAGE Keyword Syntax

ALL

Includes all data sets in the analysis.

ACTIVE (ACTV)

Includes only active data sets. Accessed as read or as write and read, these data sets are not HSM, TEMP, BACKUP, SINGLE, or BCOPY.

BACKUP (BKUP)

Includes data sets that have been opened and processed for output only and are not HSM, TEMP, SINGLE, BCOPY, or ACTIVE. Unlike active data sets, these data sets are probably being placed on tape for backup purposes.

BCOPY (BCOP)

Includes only data sets that have been opened once for input and once for output and are not HSM, TEMP, BACKUP, ACTIVE, or SINGLE. These data sets are likely a specialized form of backup. Perhaps a backup is taken and then re-copied to be sent to a vault.

TEMP

Includes only temporary data sets. Data sets are flagged in GFTAXTR as temporary if they match the system-assigned naming convention for temporary data sets or if they have been allocated as DISP=(NEW,DELETE).

SINGLE (SING)

Includes data sets that have only been accessed on a single day and are not HSM, TEMP, BCOPY, BACKUP, or ACTIVE.

HSM

Includes only those data sets owned and managed by DFSMSHsm, such as the migration level 2 and backup single-file format (SFF) data sets and autodumps. GFTAVMA has built-in knowledge of every naming convention that DFSMSHsm uses.

NONHSM

Includes all data sets except those that fall into the HSM category.

You can specify one of the following:

- **ALL**
- Any combination of **ACTIVE**, **BACKUP**, **BCOPY**, **TEMP**, and **SINGLE**
- **HSM**
- **NONHSM**

Default

If you do not specify the **USAGE** keyword, GFTAVMA will use the default, **ALL**.

VOLUME Keyword

Use **VOLUME** to include or exclude volumes from the GFTAVMA analysis based on the volume serial numbers in the input SMF records. See “Secondary Filters” on page 45 for more information.

Abbreviation: VOL

Syntax

Figure 35 shows the syntax for the **VOLUME** keyword.

[**VOLUME** (**INCLUDE** (*volume filter list*),
EXCLUDE (*volume filter list*))]

Figure 35. **VOLUME** Keyword Syntax

INCLUDE (*volume filter list*)

Specifies a volume filter list to include in the analysis.

EXCLUDE (*volume filter list*)

Specifies a volume filter list to exclude from the analysis.

See “Include/Exclude Filter Syntax” on page 45 for the rules and syntax for these keywords.

Default

If you do not specify the **VOLUME** keyword, GFTAVMA will analyze all volumes.

Relationship to Other Keywords

When filtering with **VOLUME** and **UNITADDR**, ensure that the filtering combination works. For example, if data set AA is a multivolume data set that resides on volume 1, volume 2, and volume 3, it is possible that when you use **VOLUME** and **UNITADDR** filtering, you filter volume 2 out of the analysis. This would affect the volume mount analyzer report results for data set AA.

To avoid this situation, only use these keywords to capture a totally self-contained set of volumes, such as DFSMSHsm-owned volumes. These self-contained sets of volumes have a naming convention that identifies them to be used by one particular application only.

VOLUME

Restrictions

Including or excluding specific volume serials can sometimes produce inaccurate results. For example, if data set X.Y.Z is multivolume and spans volumes 1, 2 and 3, and volume 2 gets excluded due to filtering, the data associated with data set X.Y.Z becomes meaningless.

VOLUME Example

If you specify:

```
VOLUME (EXCLUDE (B*))
```

GFTAVMA will include all volumes except those whose volume serial numbers start with B.

Chapter 6. Producing Detailed Analysis Reports

GFTAVMA produces both summary reports and detailed reports. As pointed out in “Interpreting GFTAVMA Summary Reports” on page 51, all GFTAVMA runs automatically generate summary reports. But you must use the **REPORT** keyword to generate detailed analysis reports. See “REPORT Keyword” on page 148 for keyword rules and syntax.

This chapter introduces detailed analysis reports and describes their production phases and types. It also illustrates how to interpret the actual and simulation reports.

Note: You can read GFTAVMA reports at your terminal. If you wish to print them, you should use necessary keywords to filter your input and tailor your output. Otherwise, you will generate extremely large reports that are hard to read and manage. See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 and Chapter 7, “Using GFTAVMA Keywords to Tailor Your Output,” on page 131 for more information about using these keywords.

Understanding Detailed Analysis Reports

GFTAVMA generates optional detailed reports during the three phases of GFTAVMA processing. You can use the **REPORT** keyword to request these reports. As a result, you get one output data set containing the three summary reports, described in “Interpreting GFTAVMA Summary Reports” on page 51, and any optional detailed reports you requested. This data set is defined by the SYSPRINT DD statement in the GFTAVMA JCL.

Production Phases of Detailed Reports

There are three GFTAVMA processing phases. The reports generated during each phase appear in the output data set in the order GFTAVMA generates them:

- During phase 1, GFTAVMA generates the GFTASRT1 summary report and, if requested, the Usage Report.
- During phase 2, GFTAVMA generates the GFTASRT2 summary report and, if requested, the Volume Report.
- During phase 3, GFTAVMA generates the GFTASRT3 summary report and, if requested, any or all other optional reports.

For example, if you request the Volume Report with **REPORT(VOLUME)**, the reports will be generated in the order of GFTASRT1, GFTASRT2, the Volume Report, and GFTASRT3.

Types of Detailed Reports

GFTAVMA generates two types of optional detailed reports:

- *Actual reports* that show the actual tape usage at your installation. The actual reports include:
 - The Data Set Report
 - The Maximum Gigabyte Report
 - The Top Report
 - The Usage Report

- The Volume Report
- *Simulation reports* that simulate what the tape usage at your installation might be if you use tape mount management. The simulation reports include:
 - The Estimate Report
 - The IDRC Report
 - The Management Class Report

See “Interpreting the Actual Reports” in the following section and “Interpreting the Simulation Reports” on page 106 for more details.

All volume mount analyzer messages appear in the reports. See Appendix D, “Volume Mount Analyzer Messages,” on page 167 for more information.

Interpreting the Actual Reports

The actual reports show the current tape usage statistics for your installation. There are five actual reports:

- The Data Set Report (see “The Data Set Report”)
- The Maximum Gigabyte Report (see “The Maximum Gigabyte Report” on page 88)
- The Top Report (see “The Top Report” on page 92)
- The Usage Report (see “The Usage Report” on page 95)
- The Volume Report (see “The Volume Report” on page 101)

Examples of each of the actual reports are shown in this section.

The Data Set Report

The Data Set (DSN) Report is an actual report. The actual reports show the current tape usage statistics for your installation.

Use the keyword **REPORT(DATASET)** or **REP(DSN)** to generate the Data Set Report. You can specify the **CUTOFF** keyword to reduce the number of pages printed for the Data Set Report.

If you specify:

```
REPORT(DATASET) MINSIZE(400) DSORT(SIZE)
```

GFTAVMA will generate a Data Set Report that shows data sets greater than 400MB sorted by size.

Keywords that Affect the Data Set Report

The results of the GFTAVMA reports rely on the keywords you choose. Therefore, choose your keywords carefully and be aware of the defaults and effects of the keywords on these reports.

All filtering input keywords and the **SPLIT** keyword affect these reports. See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 and “SPLIT Keyword” on page 151 for the syntax and functions of these keywords.

A subset of the tailoring output keywords specified for a GFTAVMA run also affects the Data Set Report. See Chapter 7, “Using GFTAVMA Keywords to Tailor Your Output,” on page 131 for more information about the syntax and functions of these tailoring output keywords:

- BLKSIZE
- CUTOFF

- DASDDEV
- DSORT
- FORCECAP
- GDG
- RPROGRAM

See Appendix A, “A Quick Reference to Volume Mount Analyzer Keywords,” on page 157 and Appendix B, “Volume Mount Analyzer Keyword Syntax,” on page 159 for more information about GFTAVMA keywords.

Note: It is possible for a data set on a multifile volume to show up as single-file. This happens when the data set is the first and only one written to a volume during a mount. Unless this data set is accessed again during the reporting period and a new SMF record is created, the EXTRACT portion of the volume mount analyzer will not recognize it as part of a multifile volume.

Use of the Data Set Report

The detailed Data Set Report provides further information about each data set to supplement the data set statistics in the GFTAVMA summary report, GFTASRT3.

Use the Data Set Report to examine all data sets by category, such as size, application, expiration date, and account number. The Data Set Report shows you the cost of DASD tracks for each data set as well as the cumulative DASD cost for all data sets.

Use the Data Set Report to identify data set patterns. In this way, you can determine which data sets to exclude because of size. The default for **LARGE** is 600MB. Therefore, by specifying a **MINSIZE** of 400MB for the Data Set Report, you can determine if the default for **LARGE** is accurate for your data set patterns.

The Data Set Report Legend—Explanation of Symbols

The Data Set Report is always preceded by this legend, which describes various symbols and special characters in the report. Figure 36 on page 84 is the legend for the Data Set Report.

Data Set Report

THE FOLLOWING KEYWORDS WERE INPUT TO THIS RUN: **1**
REPORT(DATASET) MINSIZE(400) DSORT(SIZE)

GFTASRT3 -- DATA SET REPORT -- FIELD KEY CHART (1 OF 2)

***** HEADING = 'CUM GB XFERD' *****

```
+-----+
| CUM GB |
| XFERD  |
+-----+
GGGGG.G   (CUMULATIVE # OF GIGABYTES)
FIELD-->1  2
```

FIELD VALUES MEANING

#1----	F	- DATA SET EXCEEDS THE DEVICE CAPACITY (DUE TO MAX BLKSIZE WITH VAR. OR UNDF. RECORDS), 'FORCECAP' WAS SPECIFIED
----	X	- DATA SET EXCEEDS THE DEVICE CAPACITY (DUE TO MAX BLKSIZE WITH VAR. OR UNDF. RECORDS), 'NOFORCECAP' WAS SPECIFIED
----	Z	- UNABLE TO DETERMINE THE SIZE OF THE DATA SET DUE TO MISSING INFORMATION (E.G., LRECL) -- FORCED SIZE = ZERO 2
#2----	U	- DATA SET HAS "UN-BLOCKED" RECORDS -- THE SIZE (MB) MAY BE ARTIFICIALLY LARGE, ALSO IT CAN'T BE RE-BLOCKED
----	V	- DATA SET HAS "VARIABLE" RECORDS -- THE SIZE (MB) MAY BE ARTIFICIALLY LARGE

***** HEADING = 'SZ-MB XFERD' *****

```
+-----+
| SZ-MB  |
| XFERD  |
+-----+
MMMMM    (CUMULATIVE # OF MEGABYTES)
FIELD----->1
```

FIELD VALUES MEANING

#1----	I	- DATA SET WAS PREVIOUSLY COMPACTED USING IDRC (IMPROVED DATA RECORDING CAPABILITY) ON THE DEVICE IT CURRENTLY RESIDES ON -- THIS CAN ONLY OCCUR ON A 348X DEVICE AND ABOVE
--------	---	---

***** HEADING = '3380E TRACKS' *****

```
+-----+
| 3380E |
| TRACKS|
+-----+
TTTTTT   (# OF 3380E TRACKS) 3
FIELD--->1  2
OR
NNNNNK   (# OF 3380E K-TRACKS -- 1000S OF TRACKS)
OR
>BBBBB   (WHERE "BBBBB" = BLKSIZE OF THE DATA SET 4)
```

Figure 36. The Data Set Report Legend (Part 1 of 2)

GFTASRT3 -- DATA SET REPORT -- FIELD KEY CHART (2 OF 2)

```

FIELD  VALUES  MEANING
-----
#1----> > - DATA SET IS CURRENTLY BLOCKED > 1/2 TRK AND EITHER "UN-BLOCKED" RECORDS OR PROGRAM
        WON'T ALLOW RE-BLOCKING
        ----> Z - THE SMF TYPE 14/15 RECORD HAD A BLKSIZE=0 AND THE DEFAULT/SPECIFIED VALUE WAS USED

#2----> S - THE 'SPACE' KEYWORD WAS ORIGINALLY CODED IN THE JCL FOR THIS DATA SET

                ***** HEADING = 'PROGRAM' *****

        +-----+
        | PROGRAM |
        +-----+
        P P P P P P P P (NAME OF THE CREATING PROGRAM -- IF KNOWN)
FIELD--->1

FIELD  VALUES  MEANING
-----
#1----> > - THE PROGRAM LISTED MAY NOT BE THE PROGRAM WHICH CREATED THIS DATA SET -- SINCE NO
        OUTPUT RECORDS WERE FOUND FOR THIS DATA SET, IT WAS IMPOSSIBLE TO DETERMINE THE ACTUAL
        CREATION PROGRAM

***** HEADING = 'DATA SET NAME' *****

        +-----+
        | DATA SET NAME |
        +-----+
        D D D D D . . . (DATA SET NAME - 43 CHARACTERS)
FIELD--->1

FIELD  VALUES  MEANING
-----
#1----> > - THE DATA SET MAY NOT BE CATALOGED -- THE GFTAVMA PROGRAM WAS UNABLE TO DETERMINE
        W/O ISSUING A LOCATE 5
        ----> @ - THE DATA SET MAY NOT BE CATALOGED AND LABEL=(,NL) WAS SPECIFIED IN THE JCL
        NOTE: IF OTHER DATA SETS WITH THIS HLQ ARE CATALOGED, THEN THIS DATA SET IS
        PROBABLY CATALOGED
        ----> + - THE DATA SET NAME LISTED IS A TEMPORARY DATA SET (BASED ON FLAGS IN THE SMF RECORD)
        -- UNCATALOGED
        ----> # - THE DATA SET NAME LISTED IS A TEMPORARY DATA SET AND LABEL=(,NL) WAS SPECIFIED
        IN THE JCL
        NOTE: NOT ALL TEMPORARY DATA SETS HAVE THE FORMAT OF 'SYSYDDDD.THHMSS.RA000.JOB.*'

```

Figure 36. The Data Set Report Legend (Part 2 of 2)

1 Shows the keywords that were specified in the GFTAVMA JCL. This information shows what keywords generated the Data Set Report.

2 See the **FORCECAP** keyword in “FORCECAP/NOFORCECAP Keywords” on page 140 for a discussion of what happens when the data set size exceeds the physical size of the tape cartridge. Also see “Tape Cartridge Usage” on page 35 that explains the capacities of tape volumes.

3 The heading and actual track calculations for this sample report are based on the **DASDDEV** keyword, which defaulted to 3390-3.

Data Set Report

4 The data set block size value affects how efficiently DASD track space is allocated. GFTAVMA calculates the DASD track size using the current data set block size.

5 "W/O ISSUING A LOCATE" means "without reading the catalog". See note 8.

Note: This block size might be efficient for tape. However, it is inefficient for DASD track usage.

Example of the Data Set Report

Note: Figure 37 on page 87 is only *part* of the complete Data Set Report that GFTAVMA generated for this run. These data sets are between 402 MB and 499 MB in size.

GFTASRT3 -- DATA SET REPORT

EXPDT	CUM STAT MOUNTS	STAT MOUNTS	CUM GB XFERD	SZ-MB XFERD	3380E TRACKS	#INP OPNS	#OUT OPNS	# GDS	DSN#	JOBNAME	PROGRAM	DATA SET NAME
04003	2.0	2.0	0.4V	402I	8486	0	2	1	1	AWRB999D	ABCDE01	AWRB.DEPRDS.BACKUP.G1150V00
04011	3.5	1.5	0.8	402I	8551	0	2	1	2	AWRB228D	DFSUDMP0	AWRB.LIB.DPTRNAPR.H2914V00
04011	5.0	1.5	1.2	402I	8551	0	2	1	3	AWRB400M	DFSUDMP0	AWRB.LIB.DPTRNAPR.H2915V00
04011	6.5	1.5	1.6	402I	8551	0	2	1	4	AWRB228D	DFSUDMP0	AWRB.OFFSITE.DPTRNAPR.H2915V00
04011	8.0	1.5	2.0	402I	8551	0	2	1	5	AWRB400M	DFSUDMP0	AWRB.OFFSITE.DPTRNAPR.H2916V00
	8.5	0.5	2.4V	410I	8735	1	0		6	C991DS01	>IEBGENER	>DEPCIF.M.DGBAPR30.MTRQRPGS
04002	16.5	8.0	2.8U	410I	12522	0	8		7	PLRC23W	XYZ	>PLRC.XYZ.WEEKLY.RESCUE.BACKUP
04002	24.5	8.0	3.2U	410I	12522	0	8		8	PLRC23W	XYZ	>PLRC.XYZ.WEEKLY.RESCUE.BACKUP.COPY
	25.5	1.0	3.7	410I	12545	3	0		9	C997MCT2	>GIMSMP	@SMPTFIN
04003	73.5	48.0	4.1V	429I	8718	24	24		10	BOFO06D	EP5260	>BOGO.ACH.BKUP.E9130
03180	75.5	2.0	4.5U	433I	8761	0	2		11	C997JMC5	IEBCOPY	>C997JMC.SMPPTS.BACKUP
91158	77.5	2.0	5.0V	457I	9739	0	2	1	12	AWRB200D	DFSRRC00	AWRB.LOG.H1010V00
91163	79.5	2.0	5.4V	457I	9745	0	3	1	13	AWRB200D	DFSRRC00	AWRB.LOG.H1008V00
91180	119.6	40.1	5.9V	458I	9749	0	41	1	14	AWRB200D	DFSRRC00	AWRB.LOG.H1005V00
	120.6	1.0	6.3V	462I	14109	1	0	1	15	C997RMOA	>IFASMFDP	PLRC.SMF.COMBINED.DAILY.H1897V00
91168	132.6	12.0	6.8V	467I	14606	9	6		16	PLRC804M	DFSRRC00	>PLRC.PLRC804.STEP06
91168	162.6	30.0	7.3V	467I	14609	24	6		17	PLRC804M	DFSURG10	>PLRC.PLRC804.STEP07
	171.6	9.0	7.8	478	9727	9	0		18	CTRU010R	>IEBGENER	@QWERTY.SSSS
	172.6	1.0	8.2V	480I	14657	1	0	1	19	C997RMOA	>IFASMFDP	PLRC.SMF.COMBINED.DAILY.H1894V00
04007	172.9	0.3	8.7V	486I	10360	0	2	1	20	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1269V00
04007	173.2	0.3	9.2V	486I	10360	0	2	1	21	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1268V00
04007	173.5	0.3	9.7V	488I	10403	0	2	1	22	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1268V00
04007	173.8	0.3	10.2V	488I	10403	0	2	1	23	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1267V00
04007	174.1	0.3	10.7V	491I	10452	0	2	1	24	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1289V00
04007	174.4	0.3	11.2V	491I	10452	0	2	1	25	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1288V00
04007	174.7	0.3	11.7V	493I	10491	0	2	1	26	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1279V00
04007	175.0	0.3	12.2V	493I	10491	0	2	1	27	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1278V00
04007	175.3	0.3	12.6V	493I	10499	0	2	1	28	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1277V00
04007	175.6	0.3	13.1V	493I	10499	0	2	1	29	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1276V00
04007	175.9	0.3	13.6V	493I	10503	0	2	1	30	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1278V00
04007	176.2	0.3	14.1V	493I	10503	0	2	1	31	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1277V00
04007	176.5	0.3	14.6V	493I	10507	0	2	1	32	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1283V00
04007	176.8	0.3	15.1V	493I	10507	0	2	1	33	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1282V00
04007	177.1	0.3	15.6V	493I	10509	0	2	1	34	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1274V00
04007	177.4	0.3	16.1V	493I	10509	0	2	1	35	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1273V00
	178.4	1.0	16.6V	493I	15051	1	0	1	36	C997RMOA	>IFASMFDP	PLRC.SMF.COMBINED.DAILY.H1898V00
04007	178.7	0.3	17.1V	496I	10574	0	2	1	37	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1282V00
04007	178.9	0.3	17.6V	496I	10574	0	2	1	38	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1281V00
	180.9	2.0	18.1V	497I	15157	2	0	1	39	C997RMOA	>IFASMFDP	PLRC.SMF.COMBINED.DAILY.H1893V00
04003	189.9	9.0	18.6U	497I	15163	0	9		40	PLRC126W	XYZ	>PLRC.XYZ.PROD14.BACKUP
05002	198.9	9.0	19.1U	497I	15163	0	9		41	PLRC126W	XYZ	>PLRC.XYZ.PROD14.BACKUP.COPY
05007	199.2	0.3	19.6V	498I	10604	0	2	1	42	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1284V00
05007	199.5	0.3	20.1V	498I	10604	0	2	1	43	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1283V00
05007	199.8	0.3	20.6V	498I	10610	0	2	1	44	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1287V00
05007	200.1	0.3	21.1V	498I	10610	0	2	1	45	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1286V00
04007	200.4	0.3	21.6V	499I	10621	0	2	1	46	AWRB250D	DPCOPY02	AWRB.LIB.MTRQRPGS.H1275V00
04007	200.7	0.3	22.1V	499I	10621	0	2	1	47	AWRB250D	DPCOPY02	AWRB.OFFSITE.MTRQRPGS.H1274V00

GFTAVMA PROGRAM COMPLETED AT 10:06:42 A.M. ON 05/22/2004 -- RETURN CODE = 0

Figure 37. Partial Data Set Report

6 Number of statistical mounts attributed to each data set. See "Statistical Mounts" on page 30 for more information.

Data Set Report

7 The size of the data set the last time it was written. If the data set was not written in the input sample, this is the size of the data transfer during the last read operation.

If you specified **GDG(GROUP)**, the SZ-MB XFERD and TRACKS size fields are the cumulative total size of each GDS within the GDG.

8 The amount of DASD space that would be required to hold the data set. If the block size is less than a half track of the device, the block size is shown. The target DASD device type used here is controlled by the **DASDDEV** keyword, which defaulted to 3380E in this report example.

9 Reading a row of the Data Set Report shows that the latest version of this data set is named PLRC.XYZ.PROD14.BACKUP.COPY. The > symbol directly in front of the data set name indicates that the data set might not be cataloged. The volume mount analyzer determines catalog status from the data set JFCB, which is contained in the SMF 14 and 15 record types. The volume mount analyzer does not have access to the catalog.

This data set was created by the job named PLRC126W and by the program named XYZ. It has an expiration date (EXPDT) of 99002. This expiration date indicates that a tape management program controls the data set.

During the reporting period, GFTAVMA determined that there were nine statistical mounts that were attributed to the data set. See “Statistical Mounts” on page 30 for more details.

The U in the CUM GB XFERD column indicates that the data set has RECFM=U. When it was created, the size of the data transfer for this data set was 497MB. The I to the right of 497 in the SZ-MB XFERD column indicates that the data set has been compacted by IDRC.

This data set was opened for output nine times, which is equal to the statistical mounts value, and never opened for input. This means that this data set is backup. See “Data Set Usage Types” on page 36 for details.

10 The number of OPENS for INPUT and OPENS for OUTPUT for a data set. These numbers show the usage pattern of active versus backup, assuming the data set is not marked USAGE(SINGLE) or USAGE(BCOPY).

11 If the data set is in a GDG, this column shows the number of individual GDSs it represents. This number is only greater than 1 if you specify **GDG(GROUP)**.

12 Program and job that created the data set or last read if GFTAVMA cannot determine the creator.

The Maximum Gigabyte Report

The Maximum Gigabyte (GBMAX) Report is an actual report. The actual reports show the current tape usage statistics for your installation.

Use the keyword **REPORT(GBMAX)** or **REP(GB)** to generate the GBMAX Report.

If you specify:

```
REPORT(GBMAX)
```


GFTAVMA will generate a Maximum Gigabyte Report.

Keywords that Affect the Maximum Gigabyte Report

The results of the GFTAVMA reports rely on the keywords you choose. Therefore, choose your keywords carefully and be aware of the defaults and effects of the keywords on these reports.

All filtering input keywords and the **SPLIT** keyword affect these reports. See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 and “SPLIT Keyword” on page 151 for the syntax and functions of these keywords.

A subset of the tailoring output keywords specified for a GFTAVMA run affect the GBMAX Report. See Chapter 7, “Using GFTAVMA Keywords to Tailor Your Output,” on page 131 for more information about the syntax and functions of these tailoring output keywords:

- BLKSIZE
- FORCECAP
- RPROGRAM

See Appendix A, “A Quick Reference to Volume Mount Analyzer Keywords,” on page 157 and Appendix B, “Volume Mount Analyzer Keyword Syntax,” on page 159 for more information about GFTAVMA keywords.

Use of the Maximum Gigabyte Report

The Maximum Gigabyte Report shows the number of gigabytes for the DISP=NEW and DISP=MOD data sets being created for every hour of every day in the sample.

If your volume mount analyzer run includes *all* of the data you plan on intercepting with ACS filters for tape mount management, the GBMAX Report is an excellent estimate of the amount of free space you will need in the tape mount management DASD buffer.

You can determine the amount of free space you need by first running the Estimate and Management Class Reports for each of your tape mount management target applications. Based on these reports, you might determine that you need twenty 3390-3 volumes for primary storage. To determine how much free space you need, assuming DFSMSshsm migrates data sets hourly, you would run the GBMAX Report on your final set of filters.

This example report shows that you need two 3390-3 volumes to handle your peak loads. So you would need to start with 22 primary storage volumes to allow two volumes of free space (10%) for the peak loads.

You can also use the GBMAX Report with the Estimate Report to determine the amount of buffering needed for every hour of each day. Keep in mind, however, that the maximum number of gigabytes for all applications is not the sum of the maximum for each application. Applications with maximums of 3, 4, and 5GB might only need 5GB when mixed, not 12GB.

GFTAVMA calculates the numbers in the Estimate Report (see “The Estimate Report” on page 106) by using the LRECL, RECFM, and BLKSIZE of the data sets and the number of blocks as they would lay out on the target DASD device. Therefore, do not expect the number of gigabytes in the Estimate Report to match the numbers in the GBMAX Report.

GBMAX Report

See *MVS/ESA SML: Managing Data* and *z/OS DFSMS: Implementing System-Managed Storage* for more information on using the GBMAX and Estimate Reports together.

Example of the Maximum Gigabyte Report

The GBMAX Report shows the amount of new data transferred over time for the data sample analyzed in the GFTAVMA run. The report is meant to give a perspective of how much data per hour will be allocated to a tape mount management DASD buffer if DFSMS intercepted this set of data and redirected it to the tape mount management DASD buffer.

Figure 38 on page 91 is an example of the GBMAX Report.

																								1		2	
MM/DD/YYYY	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	MAX	TOTAL	
05/27/1991	0.1	3	2	0	<.1	0.9	0.6	1	0.3	0.6	0	0	0.5	0.9	0	0.1	0.7	0.3	<.1	2	6	0.7	0.8	0	6	20	
05/28/1991	4	0.4	2	0	1	2	5	0.6	1	0.9	1	0.6	1	0.6	0.3	1	1	1	2	8	8	4	2	3	8	51	
05/29/1991	4	5	4	1	4	5	3	0.1	3	2	0.3	0.2	0.1	0.2	0.3	0.3	2	2	3	8	7	3	1	6	8	62	
05/30/1991	5	3	3	0.9	0.5	6	1	0.1	2	1	0.4	0.3	0.5	0.4	0.3	0.9	0.9	1	3	2	9	6	2	5	9	55	
05/31/1991	5	4	3	1	4	6	1	0.1	2	0.7	0.6	0.3	0.3	0.4	0.4	0.7	2	0.3	1	4	10	4	3	5	10	60	
06/01/1991	2	4	6	5	5	2	2	5	16	18	14	15	10	12	1	<.1	0.2	0.9	0.2	3	<.1	0.5	1	2	18	125	
06/02/1991	<.1	0	8	9	2	0.2	2	<.1	<.1	0	0	0	0	<.1	0	0	0	0	0	0	0	0	0.9	3	0	9	25
06/03/1991	0.1	3	3	0.4	4	0.5	3	5	2	0.9	0.4	0.4	0.6	0.4	2	0.8	1	1	2	8	8	3	1	4	8	54	
06/04/1991	0.3	2	4	5	3	3	8	<.1	2	0.6	0.3	0.3	0.5	0.7	0.2	1	1	1	2	6	8	3	2	4	8	57	
06/05/1991	1	0.5	7	5	4	4	5	3	4	3	0.2	0.1	0.5	0.1	0.7	0.7	1	3	3	6	8	4	0.8	4	8	70	
06/06/1991	2	6	5	1	1	4	4	1	2	2	2	0.5	1	0.5	0.6	0.9	1	2	2	2	10	5	2	6	10	64	
06/07/1991	3	4	6	4	4	4	2	0.5	0.9	6	0.2	0.2	<.1	0.8	0.2	1	1	0.8	2	3	10	4	3	4	10	67	
06/08/1991	1	5	6	3	4	0.6	0.9	7	18	15	17	18	13	3	0.2	0	0.2	0	3	4	4	2	4	0	18	130	
06/09/1991	0.2	4	7	5	3	1	2	0.9	0.3	0	0	0	<.1	0.6	0.6	0	0	0	0	0.3	0.9	<.1	3	0	7	28	
06/10/1991	0.1	3	2	1	4	0.3	3	4	2	3	0.7	0.3	0.3	0	2	0.9	1	2	3	9	4	4	1	3	9	54	
06/11/1991	1	4	6	3	2	4	4	0.4	2	0.7	0.2	0.2	0.3	0.4	0.1	1	2	2	2	6	9	2	2	5	9	59	
06/12/1991	4	4	4	3	3	4	3	0.3	2	0.5	0.4	<.1	0.7	1	1	1	0.8	0.9	1	7	8	3	0.8	4	8	59	
06/13/1991	2	3	6	4	1	5	4	0.4	2	0.3	0.4	<.1	0.2	0.5	0.2	0.6	1	2	3	1	9	6	1	6	9	57	
06/14/1991	2	3	6	3	3	5	4	0.1	3	0.8	0.4	0.1	0.2	0.7	0.3	0.1	1	0.9	3	3	9	5	0.9	6	9	60	
06/15/1991	3	1	3	4	5	3	4	15	16	17	18	15	6	0.6	0.7	0	0.2	0	<.1	1	<.1	0.2	0.3	1	18	114	
\																											
/																											
\																											
/																											
\																											
/																											
06/22/1991	2	5	6	2	4	0.6	7	18	16	17	18	15	0.9	0	0	0	0.2	0	0	4	<.1	0	4	0.8	18	120	
06/23/1991	0	0	0	0.6	0.6	9	5	0	<.1	0	0	0	0	0	<.1	0	0	0	0	0	0	<.1	3	0	9	19	
06/24/1991	<.1	3	2	0.3	1	4	0.1	0.3	2	2	5	1	0.1	0.5	0.8	0.7	0.8	1	2	8	5	3	1	4	8	51	
06/25/1991	1	0.1	7	6	0.5	2	5	4	3	5	4	0.3	0.6	0.7	0.3	0.9	2	2	2	7	6	3	2	5	7	67	
06/26/1991	5	3	4	2	3	5	2	0.3	1	0.8	0.2	0.2	0.2	0.7	0.6	0.7	2	1	2	7	10	3	1	4	10	59	
06/27/1991	3	4	4	2	0.9	5	2	<.1	1	0.2	1	0.2	0.3	0.3	0.4	0.7	2	3	3	3	9	6	2	5	9	60	
06/28/1991	4	3	2	3	3	3	7	3	1	2	0.1	2	0.5	0.2	0.6	1	1	0.8	2	4	8	5	5	4	8	64	
\																											
/																											
\																											
/																											

MAXIMUM GIGABYTE ALLOCATIONS BY HOUR

																								3		4	
MAX GB/HR	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	MAX	TOTAL	
HOURLY MAX	5	6	9	9	5	9	8	18	18	18	18	18	13	12	3	4	3	3	4	9	10	6	5	6	18	2060	

GFTAVMA PROGRAM COMPLETED AT 11:27:39 A.M. ON 06/03/1992 -- RETURN CODE

Figure 38. The GBMAX Report

GBMAX Report

Reading across the rows of the report by date shows the number of gigabyte of DISP=NEW and DISP=MOD data sets transferred per hour, the highest number of gigabytes transferred during the peak hour, and the total number of gigabytes transferred for that day.

Reading down a column of the report shows the highest number of gigabytes transferred during that hour of each day for the duration of the study.

The number in the HOURLY MAX row shows the highest number of gigabytes transferred during that hour (3GB was the peak for hour 17, and 18GB was the peak for hour 10) for the duration of the study. It is not the total transferred.

1 The MAX per day shows the maximum number of gigabytes for DISP=NEW and DISP=MOD data sets allocated during any one hour interval on a given day.

2 The TOTAL column shows the total number of gigabytes for DISP=NEW and DISP=MOD data sets allocated by day.

3 The MAX in the HOURLY MAX row represents the maximum number of gigabytes for DISP=NEW and DISP=MOD data sets allocated in any one hour interval throughout the sample. The MAX of 18 occurred in hours 7, 8, 9, 10, and 11. The MAX value is not a total.

Note: This MAX is the number that you should look at to determine the peak number of gigabytes for the run, which is the size needed for the hourly free space in tape mount management DASD buffer.

4 In this example, 2060GB are transferred for the duration of your study. The largest number of gigabytes transferred during a single hour for the duration of your study is 18GB. Based on this information, you would know that you need a tape mount management DASD buffer that can handle at least 18GB of data, which is the number of gigabytes transferred during the peak periods of data transfer.

Based on the report in this example, to handle the maximum amount of free space needed each hour, you would need 18GB. To buffer this set of data for 24 hours, you would need to handle the maximum of 125GB on 06/01/2003. GFTAVMA calculates the numbers in the GBMAX Report by multiplying the number of blocks by the maximum block size and then rounding up to the nearest gigabyte.

If the tape mount management DASD buffer is not migrated hourly, you should plan on a buffer that can handle at least the maximum daily tape usage of 130GB, which occurred on 06/08/91.

See *MVS/ESA SML: Managing Storage Groups* and *z/OS DFSMS: Implementing System-Managed Storage* for details on planning your DASD buffer.

The Top Report

The Top Report is an actual report. The actual reports show the current tape usage statistics for your installation. The Top Report shows the current tape usage statistics in descending order of mounts for various data set-related categories.

Use the keyword **REPORT(TOP)** or **REP(TOP)** to generate the Top Report.

In this example, these keywords were specified:

```
REPORT(TOP(PROGRAM,PERCENT(75)))
```

These keywords specify that GFTAVMA generates a Top Report to show the programs responsible for the highest number of tape mounts at your installation. PERCENT (75) is the cutoff point. Therefore, the Top Report example only shows the top 75% of cumulative tape mounts.

See “Parameters for the Top Report” on page 149 for the rules, syntax, and other kinds of Top Reports you can request.

Keywords that Affect the Top Report

The results of the GFTAVMA reports rely on the keywords you choose. Therefore, choose your keywords carefully and be aware of the defaults and effects of the keywords on these reports.

All filtering input keywords and the **SPLIT** keyword affect these reports. See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 and “SPLIT Keyword” on page 151 for the syntax and functions of these keywords.

A subset of the tailoring output keywords specified for a GFTAVMA run also affects the Top Report. See Chapter 7, “Using GFTAVMA Keywords to Tailor Your Output,” on page 131, “Parameters for the Top Report” on page 149, and “ACCTFLD Keyword” on page 15 for more information about the syntax and functions of these tailoring output keywords:

- ACCTFLD
- BLKSIZE
- TOP parameters

See Appendix A, “A Quick Reference to Volume Mount Analyzer Keywords,” on page 157 and Appendix B, “Volume Mount Analyzer Keyword Syntax,” on page 159 for more information about GFTAVMA keywords.

Use of the Top Report

The purpose of the Top Report is to show how you can use simple filters with ACS routines to intercept large numbers of data sets. In the example shown in Figure 39 on page 94, only four programs are responsible for approximately 50 percent of all tape mounts.

You can use this information to filter on these four program names to intercept 21771 data sets. Most of these data sets are probably forms of backup and could be migrated directly to migration level 2. This means that a small amount of DASD would be required to hold the data for a short period of time. In this way, all but extremely large data sets could be intercepted. It would be easier to **EXCLUDE** the exceptions by program name and keep your filter list simple than to list all the data sets to **INCLUDE** by their name.

If you want to intercept data sets that are being allocated to tape and stack them, **PROGRAM** and **UNIT=TAPE** would be good ACS filters to direct these data sets to the tape mount management DASD buffer before allocating them to tape.

The Top Report also shows the average size of the data sets and the number of data sets greater than the **LARGE** keyword value for each category. GFTAVMA assumes that you have set the **LARGE** keyword to a value that represents the largest size of data sets that you want to intercept and send to the tape mount management DASD buffer. See “LARGE Keyword” on page 144.

You should list data sets larger than the **LARGE** value in an **EXCLUDE** list and allow them to be written directly to tape. If there are few data sets in the category

Top Report

greater than **LARGE**, they could probably easily be accommodated on the tape mount management DASD buffer for a short period of time.

Example of the Top Report

Figure 39 shows the programs that are responsible for the top 75% of the tape mounts in this GFTAVMA run.

THE FOLLOWING KEYWORDS WERE INPUT TO THIS RUN:
REPORT(TOP(PROGRAM,PERCENT(75)))

DFTASRT3 -- REPORT TOP PROGRAM NAMES -- 75% MOUNTS -- "LARGE" = 600 MB

RANK	PROGRAM	# DSNs	% TOT	CUM DSN	% TOT	< LARGE	AVG SIZE	# MOUNTS	% TOT	CUM MNTS	% TOT	SIZE (GB)	% TOT	CUM GB	% TOT
1	IDCAMS	12386	30.0	12386	30.0	0	16.0	6611.6	17.3	6611.6	17.3	198.3	15.1	198.3	15.1
2	IEBGENER	3273	7.9	15659	38.0	3	7.9	5429.7	14.2	12041.3	31.4	25.8	2.0	224.1	17.1
3	DFSUDMPO	5966	14.5	21625	52.4	0	51.4	4223.9	11.0	16265.3	42.5	306.8	23.4	530.8	40.5
4	XYZ	146	0.4	21771	52.8	130	837.5	2673.0	7.0	18938.3	49.4	122.3	9.3	653.1	49.9
5	ABCDE04	5218	12.7	26989	65.4	0	6.5	2447.9	6.4	21386.2	55.8	34.1	2.6	687.2	52.5
6	ABCDE01	1320	3.2	28309	68.6	2	23.9	1617.0	4.2	23003.2	60.1	31.5	2.4	718.8	54.9
7	IFASMFDP	69	0.2	28378	68.8	10	358.4	1333.0	3.5	24336.2	63.5	24.7	1.9	743.5	56.8
8	DFSRR00	636	1.5	29014	70.4	14	81.0	1229.1	3.2	25565.3	66.7	51.5	3.9	795.0	60.7
9	DUMP	477	1.2	29491	71.5	0	11.6	955.7	2.5	26521.0	69.2	5.5	0.4	800.5	61.1
10	TATTGET4	92	0.2	29583	71.7	87	1102.9	819.0	2.1	27340.0	71.4	101.5	7.7	902.0	68.9
11	LMNOPR	209	0.5	29792	72.2	0	31.1	708.0	1.8	28048.0	73.2	6.5	0.5	908.4	69.4
12	DKNMTASK	140	0.3	29932	72.6	0	8.7	523.0	1.4	28571.0	74.6	1.2	0.1	909.7	69.5
13	PANNN	92	0.2	30024	72.8	0	107.0	486.0	1.3	29057.0	75.9	9.8	0.8	919.5	70.2

GFTAVMA PROGRAM COMPLETED AT 10:08:27 A.M. ON 06/03/1992 -- RETURN CODE =

Figure 39. The Top Report

1 # MOUNTS % TOT column shows the number of mounts and the percent of the total mounts for each RANK. The Top Report is ranked according to the associated number of mounts sorted in descending order, and RANK 1 is the program responsible for the most number of mounts at this installation.

2 CUM MNTS % TOT column shows the total mounts for the programs. Usually less than 20 programs account for more than 70% of the mounts. In this example, 13 programs were responsible for 75.9% of the mounts.

Reading across the rows of this example, you see that the first row (RANK 1) shows that the program IDCAMS created 12386 data sets (# DSNs column), which account for 17.3% of the mounts (% TOTAL # OF MOUNTS column). These data sets have an average size of 16.0MB (AVG SIZE column). None of these data sets are greater than the **LARGE** value of 600MB (> LARGE column).

Contrast this with the second row of the report (RANK 2), which shows that 3273 data sets were created by the program IEBGENER (# DSNs column). These data sets account for 14.2% of the mounts (% TOTAL # OF MOUNTS column). The average data set size is only 7.9MB, and only three of the data sets are greater than the **LARGE** value. Therefore, you can reduce many mounts by using just one program filter (IEBGENER).

The fourth row (RANK 4) shows that only four programs are responsible for almost half (49.4%) of all mounts (% TOTAL CUM MNTS column). These same

four programs are also responsible for creating 52.8% of all the tape data sets in this run (% TOTAL OF CUM DSN column).

Therefore, you can filter on four programs and reduce your mounts by 49.4%.

Note: The Top Report can be used to focus on several categories, such as program name, job name, high level qualifier, or account number. See “Parameters for the Top Report” on page 149 for the rules, syntax, and other types of Top Reports you can request.

The Usage Report

The Usage Report is an actual report. The actual reports show the current tape usage statistics for your installation.

Use the keyword **REPORT(USAGE)** or **REP(USE)** to generate the Usage Report.

If you specify:

REPORT(USAGE)

GFTAVMA will generate a Usage Report.

Keywords that Affect the Usage Report

All filtering input keywords and the **SPLIT** keyword affect the Usage Report. See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 and “SPLIT Keyword” on page 151.

None of the tailoring output keywords affects the Usage Report.

Description

The Usage Report shows the breakdown of tape mount activity and tape drive usage by one hour intervals. Since this report is generated *after* filtering, the results are for filtered data.

The Usage Report does not show the effects of IDRC compaction, 36 track, or Enhanced Capacity Cartridge System Tape.

You should be aware of shared system considerations when running the Usage Report. See “Understanding the Limitations of the Volume Mount Analyzer” on page 5.

GFTAVMA Ways of Counting Mounts and Allocations

GFTAVMA processes both demount (Type 21) records and open (Types 14 and 15) records to obtain the mount and allocation information. Since the mount time is not represented in either record, GFTAXTR matches the volume demount with one or more data set opens to determine when and how long the volume was actually opened. As a result, the number of transports the volume mount analyzer determines to be in use might be slightly lower (10%) than the number actually in use.

The Tape Mounts per Hour section of the Usage Report is a count of all mounts that occurs during each hourly interval. The Maximum Concurrent Tape Allocations per Hour section of the report is based on exactly the same set of data, but the data is examined at 10 second intervals to determine the maximum number of tapes that were ever concurrently mounted. The report shows the maximum, not the total, for each hour.

The Usage Report does not provide absolute precision because the time between a demount and allocation can be less than 10 seconds and GFTAVMA rounds to 10 second intervals. See Figure 40 on page 97.

New Technology Tape Drives That You Need

The Usage Report analyzes tape activity that has already occurred; it does not project future activity. Determining the number of tape drives of newer device types needed in the future is dependent upon several assumptions, many of which GFTAVMA cannot make. For instance:

- If IDRC is used to compact data at a 3:1 compaction ratio, how much time will the job need a tape drive?
- How much will processing time decrease due to improvements in hardware technology in the tape subsystem (increased buffer size, faster channel speed, decreased rewind time)?
- If times for job processing are shortened, how will GFTAVMA predict the changes in the processing schedule so as to be able to determine the number of tape drives that would be used concurrently?
- If DFSMS is used to manage tape data sets in the storage hierarchy and to place data sets on migration level 2 tape, at what times should it be assumed that DFSMSHsm Space Management will run?

If tape mounts are reduced by 50-70% and managed entirely by DFSMSHsm, the reduction will be so dramatic that it is practically impossible to emulate on current technology and certainly impossible to emulate on newer technology. The volume mount analyzer reports, however, can be useful in monitoring the effectiveness of your new technology. For example, you might determine how many 348X drives you can eliminate after implementing 3490s.

Use of the Usage Report

Use the Usage Report to determine how many tape mounts you can eliminate using tape mount management. You can also get some indication of what impact tape mount management might have on reducing your batch window.

If you derive a set of ACS filters to capture a particular data set for DFSMS to intercept and manage, this report can show the amount of mounts that will no longer occur and the times when the mounts will be reduced.

You can also study the number of tape devices that are tied up servicing this set of data and the distribution of tape usage over time. To replace this tape usage activity, you can determine the amount of data that would be written to migration level 2 and plan the additional workload that would be given to DFSMSHsm. You can also decide the times at which DFSMSHsm processes this workload.

See *z/OS DFSMS: Implementing System-Managed Storage* for an explanation of tape mount management and DFSMSHsm. See *z/OS DFSMSHsm Implementation and Customization Guide* for more details about DFSMSHsm.

The Usage Report also shows how effective your automated tape libraries are compared to your manual tape transports. To make this comparison, filter on the UNIT addresses for the automated transports and for the manual transports. If there is no clear difference between the numbers, automation is probably helping the symptom but not the cause of your problems.

Note: The volume mount analyzer cannot project hardware requirements or show you how to redistribute your workload because the number of parameters needed for the simulation will be prohibitive.

Example of the Usage Report (Tape Allocations)

Figure 40, Figure 41 on page 98, and Figure 42 on page 99 provide an example of the Usage Report. These parts of the report describe the tape allocations.

GFTASRT1 -- USAGE REPORT -- TAPE DRIVE BY SYSTEM STATISTICS

DEVTYPE	ADDR	MOUNTS	SMF SYSTEM IDENTIFIERS (SID'S)	1
3420	0280	821	SYSA	
	0281	205	SYSA	
348X	0520	3024	SYSA	
	0521	2628	SYSA	
	0522	2238	SYSA	
	0523	2396	SYSA	
	0524	2503	SYSA	
	0525	2298	SYSA	
	0526	2142	SYSA	
	0527	2220	SYSA	
	0528	2012	SYSA	
	0529	2239	SYSA	
	052A	2133	SYSA	
	052B	2229	SYSA	
	052C	2271	SYSA	
	052D	2046	SYSA	
	052E	2063	SYSA	
	052F	1784	SYSA	
3420	0880	1013	SYSA	
	0881	321	SYSA	
4	3420	TAPE DRIVES		
16	348X	TAPE DRIVES		

20	TOTAL TAPE DRIVES IN THIS SAMPLE			
38586	TOTAL TAPE MOUNTS IN THIS SAMPLE			

Figure 40. The Usage Report—Tape Allocations, System Statistics

1 The first part of the Usage Report shows the number of mounts per device by device address.

Usage Report

GFTASRT1 -- TAPE DRIVE/MOUNT USAGE REPORT

 * MAXIMUM CONCURRENT TAPE ALLOCATIONS BY HOUR *

2

3

DATE:	HOUR:																							<==== DAILY =====>		
MM/DD/YYYY	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	MAXIMUM	MOUNTS
05/27/1991	3	4	11	2	2	2	5	2	4	2	1	0	5	6	1	1	4	2	1	3	12	6	5	2	12	277
05/28/1991	4	3	10	2	4	6	5	6	7	12	5	8	7	4	6	9	8	6	5	11	14	13	10	7	14	1274
05/29/1991	10	17	18	5	7	9	5	4	9	16	9	8	7	6	4	6	7	7	7	12	14	10	15	18	18	1525
05/30/1991	12	14	10	5	3	7	6	4	11	12	14	6	6	6	5	7	7	6	10	9	14	15	15	16	16	1372
05/31/1991	13	14	9	5	7	7	5	5	8	8	6	5	5	5	8	7	6	4	3	11	12	11	9	12	14	1367
06/01/1991	15	16	16	14	17	12	11	13	13	12	12	7	9	10	5	2	3	2	4	6	10	5	9	10	17	1103
06/02/1991	3	0	16	12	8	7	6	5	2	0	0	0	0	1	0	0	1	0	0	0	2	3	8	0	16	253
06/03/1991	3	3	12	4	4	6	6	2	8	9	9	7	8	5	9	8	6	9	7	12	13	12	7	11	13	1204
06/04/1991	6	8	17	17	15	10	8	3	9	7	7	6	6	8	6	7	8	6	6	13	16	10	8	11	17	1433
06/05/1991	11	11	16	16	17	7	7	7	7	8	6	8	7	5	7	8	12	10	8	12	16	14	8	13	17	1447
06/06/1991	16	13	16	6	9	9	5	7	10	13	11	8	8	8	9	9	14	14	9	12	16	16	12	16	16	1543
06/07/1991	12	12	14	11	8	6	4	4	7	9	6	7	4	5	6	10	11	3	7	9	14	13	11	11	14	1366
06/08/1991	12	12	17	15	14	6	7	12	11	11	9	8	10	6	2	0	2	0	10	12	4	5	9	0	17	1015
06/09/1991	3	4	16	9	6	8	6	6	3	0	0	0	1	2	2	0	2	0	1	2	3	1	8	0	16	294
06/10/1991	3	5	12	6	4	5	6	5	5	11	8	6	7	4	9	6	6	7	8	11	15	11	7	7	15	1139
06/11/1991	7	17	18	13	8	6	4	5	7	11	8	6	5	4	5	8	8	7	6	10	13	12	9	13	18	1507
06/12/1991	11	14	15	6	8	6	5	6	11	8	9	9	6	9	8	7	6	4	5	14	15	14	11	11	15	1442
06/13/1991	8	15	16	15	13	9	7	4	8	9	11	8	6	9	5	8	6	8	10	12	16	15	9	18	18	1410
06/14/1991	10	8	16	10	8	8	7	4	7	8	6	8	7	14	14	11	15	13	8	10	16	11	11	15	16	1357

Figure 41. The Usage Report—Tape Allocations by Hour

2 A tape allocation is the entire period of time a tape drive is assigned to a job and therefore unavailable for other work. It includes the time when the volume was opened (for jobs that read/write several volumes through the same DD statement) to the time when the tape was demounted.

For each one hour period of each day in the sample, the number shown in this report is the peak number of tape drives that were allocated during the one hour interval, not the total for the hour.

3 The DAILY MAXIMUM column shows the highest number of concurrent tape allocations for each hour of the day (as described in **2**). For example, on May 27, 2003, during hour 20, there were 12 concurrent tape allocations. This was the highest number for one hour period on that day, and 12 was the DAILY MAXIMUM. It was not the total for the day.

GFTASRT1 -- TAPE DRIVE/MOUNT USAGE REPORT
GFTASRT1 -- TAPE DRIVE

* MAXIMUM CONCURRENT TAPE ALLOCATIONS BY HOUR *

DATE: MM/DD/YYYY	HOUR:																							<=== DAILY ===>		
	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	MAXIMUM	MOUNTS
06/15/1991	9	10	16	14	16	11	10	13	10	10	8	8	6	3	2	2	1	0	2	4	7	3	5	6	16	914
06/16/1991	8	3	16	13	4	3	15	10	9	7	5	3	3	3	2	1	3	1	1	2	2	3	11	3	16	453
06/17/1991	3	4	11	5	5	7	5	6	10	8	6	6	6	7	7	10	9	10	10	13	14	9	8	9	14	1188
06/18/1991	8	5	15	14	16	16	7	8	9	13	8	7	8	10	9	9	11	9	9	13	16	13	15	13	16	1557
06/19/1991	9	12	14	8	6	7	5	6	9	13	11	8	7	7	7	7	8	8	7	15	16	17	16	13	17	1401
06/20/1991	9	12	15	5	6	7	4	4	8	10	10	9	8	7	8	8	9	11	9	13	16	12	13	13	16	1423
06/21/1991	9	11	12	8	8	7	7	4	11	7	6	5	4	5	5	4	6	6	8	11	14	10	10	12	14	1298
06/22/1991	14	15	14	12	10	5	12	10	10	10	9	10	4	0	0	0	2	0	2	8	8	3	12	2	15	921
06/23/1991	0	0	0	2	2	16	12	2	2	0	0	0	0	0	1	1	1	0	0	0	0	1	8	0	16	202
06/24/1991	1	2	11	2	4	10	4	3	10	10	6	8	7	10	5	5	7	4	6	11	13	13	7	12	13	1131
06/25/1991	7	4	15	3	10	14	15	15	17	14	13	11	10	9	6	11	11	9	9	11	15	13	12	13	17	1659
06/26/1991	17	14	12	7	6	9	4	6	7	7	11	7	8	8	11	11	14	8	8	14	15	12	12	13	17	1371
06/27/1991	12	10	14	10	7	10	5	3	7	7	9	7	6	6	5	6	8	9	10	11	15	15	13	12	15	1400
06/28/1991	8	11	12	11	7	10	8	5	5	4	5	5	4	4	3	5	6	6	6	10	13	14	11	10	14	1340

* MAXIMUM TAPE ALLOCATIONS: 33 DAYS SUMMARY *

HOURLY MAXIMUM:	HOUR:																							<=== 33 DAYS ===>		
	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	MAXIMUM	MOUNTS
	17	17	18	17	17	16	15	15	17	16	14	11	10	14	14	11	15	14	10	15	16	17	16	18	18	38586

Figure 42. The Usage Report—Tape Allocations by Hour (continued)

4 The MOUNTS column shows the total number of mounts that occurred for each day. This number is obtained from the Tape Mounts per Hour table in Figure 43 on page 100 and Figure 44 on page 101.

5 The maximum concurrent tape allocations for each hour in the reporting period shows the highest number of concurrent tape allocations for that hour in the entire reporting period.

6 The maximum concurrent tape allocations for any one hour in this reporting period shows the maximum number of concurrent tape allocations was 18, which occurred in hours 2 and 23. This is not the total for the period.

In Figure 42, the maximum number of drives allocated in an hour was 18 even though there were 20 drives total. The reason for this difference might be because of the difference between open time and mount time. If the data on the 3420s had to be kept on tape, you would run another Usage Report and **EXCLUDE** 3420s, so you could see a more accurate view of how your 348X drives are being used.

Note: The mount time does not always represent allocation time. In some cases, such as with JES3 systems, drives are allocated well ahead of when they are actually used. GFTAVMA does not have access to that information, so there

Usage Report

might be up to 10% variance between the actual number of drives allocated and the number of GFTAVMA reports.

Example of the Usage Report (Tape Mounts)

Figure 43 and Figure 44 on page 101 show the part of the Usage Report that describes tape mount statistics.

```

GFTASRT1 -- TAPE DRIVE/MOUNT USAGE REPORT
*****
* TAPE MOUNTS PER HOUR * 7
*****

```

DATE: MM/DD/YYYY	HOUR:																							<==== DAILY =====>		
	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	MAXIMUM	MOUNTS
05/27/1991	6	14	28	3	3	6	16	4	11	5	1	0	12	13	1	2	9	2	1	11	69	22	36	2	69	277
05/28/1991	22	10	44	2	11	10	19	22	30	75	54	59	66	30	51	96	76	60	72	98	145	107	71	44	145	1274
05/29/1991	90	107	68	45	22	22	10	18	64	49	41	60	49	48	19	35	58	52	82	146	113	102	91	134	146	1525
05/30/1991	90	41	48	18	13	14	12	20	75	63	69	29	33	41	25	47	40	60	69	58	161	115	115	116	161	1372
05/31/1991	110	50	51	19	25	18	34	22	60	61	36	20	32	17	41	66	45	16	40	118	157	89	113	127	157	1367
06/01/1991	84	124	103	55	64	48	43	54	79	80	66	54	45	54	3	1	4	6	15	38	13	11	41	18	124	1103
06/02/1991	4	0	87	50	24	9	15	3	2	0	0	0	0	1	0	0	3	0	0	0	2	8	45	0	87	253
06/03/1991	4	19	42	15	14	8	27	8	63	49	66	80	32	21	49	47	59	51	79	139	124	85	55	68	139	1204
06/04/1991	37	53	123	62	82	29	27	20	54	56	41	48	39	41	18	60	40	48	70	100	125	94	71	95	125	1433
06/05/1991	61	52	106	58	46	26	20	37	42	52	30	69	41	17	22	59	61	63	72	85	150	105	65	108	150	1447
06/06/1991	113	71	64	31	23	18	23	48	62	62	78	48	25	37	19	57	67	70	72	41	149	130	82	153	153	1543
06/07/1991	100	63	64	43	36	11	10	28	43	73	74	32	16	21	19	51	53	15	70	98	148	117	91	89	148	1366
06/08/1991	73	76	114	65	49	19	21	53	78	76	67	73	52	16	0	0	2	0	48	56	13	10	54	0	114	1015
06/09/1991	6	20	86	32	26	18	16	9	4	0	0	0	1	6	4	0	3	0	2	5	8	2	46	0	86	294
06/10/1991	4	19	37	19	10	6	30	12	26	71	45	69	36	34	51	64	50	63	83	107	118	85	52	48	118	1139
06/11/1991	64	109	90	51	32	14	15	24	66	47	57	48	32	35	27	61	59	61	69	98	136	96	89	127	136	1507
06/12/1991	90	83	65	30	23	12	25	20	88	63	56	18	38	55	33	62	43	14	75	111	112	117	74	135	135	1442
06/13/1991	49	76	78	58	34	28	15	18	45	46	71	41	21	35	16	63	36	65	72	44	149	117	88	145	149	1410
06/14/1991	64	59	67	44	22	33	22	18	56	69	37	38	30	30	42	48	53	22	83	78	121	72	83	165	165	1357

Figure 43. The Usage Report—Tape Mounts

7 The value for Tape Mounts per Hour is the cumulative number of mounts occurring on all tape drives during that hour.

8 The DAILY MAXIMUM column shows the maximum number of mounts for the peak hour of that day, and the DAILY MOUNTS column is the total number of mounts for the day.

GFTASRT1 -- TAPE DRIVE/MOUNT USAGE REPORT

```
*****
* TAPE MOUNTS PER HOUR
*****
```

DATE:	HOUR:																							<=== DAILY ===>		
MM/DD/YYYY	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	MAXIMUM	MOUNTS
06/15/1991	101	66	55	81	43	41	40	81	80	74	68	59	25	6	4	2	2	0	10	28	9	3	22	13	101	914
06/16/1991	23	7	102	47	7	1	77	22	37	41	4	4	5	1	2	0	3	0	0	1	0	7	54	8	102	453
06/17/1991	10	9	56	19	26	14	13	27	40	58	44	44	28	37	35	44	55	83	93	119	101	90	64	79	119	1188
06/18/1991	63	38	101	68	60	35	20	27	53	74	69	25	43	32	29	51	55	64	78	84	153	98	94	143	153	1557
06/19/1991	72	49	58	28	16	21	11	26	55	71	60	27	15	21	27	46	61	68	65	93	154	118	142	97	154	1401
06/20/1991	74	41	79	13	17	14	9	19	41	69	58	43	41	48	34	52	54	77	70	54	165	93	123	135	165	1423
06/21/1991	68	51	61	28	23	21	44	17	65	51	58	32	11	29	22	32	52	55	84	70	159	100	69	96	159	1298
06/22/1991	128	71	93	40	46	21	52	78	75	73	65	59	6	0	0	0	2	0	9	34	9	2	54	4	128	921
06/23/1991	0	0	0	6	4	91	35	10	5	0	0	0	0	0	1	0	3	0	0	0	0	2	45	0	91	202
06/24/1991	1	13	40	5	9	14	6	15	67	51	42	64	33	40	44	44	51	36	79	132	123	114	49	59	132	1131
06/25/1991	49	5	53	6	25	70	83	52	81	104	89	50	51	55	27	57	87	64	62	81	176	89	113	130	176	1659
06/26/1991	90	50	56	20	24	19	10	28	66	44	55	32	39	42	41	60	63	40	48	104	137	90	91	122	137	1371
06/27/1991	89	64	58	28	15	18	22	9	61	34	58	34	26	32	20	43	50	81	76	65	152	123	130	111	152	1400
06/28/1991	71	68	59	28	20	29	56	18	58	44	47	37	22	20	19	64	45	39	50	95	161	95	114	80	161	1340

```
*****
* MAXIMUM TAPE MOUNTS: 33 DAYS SUMMARY
*****
```

	HOUR:																							<=== 33 DAYS ===>			
	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16	17	18	19	20	21	22	23	MAXIMUM	MOUNTS	
HOURLY																											
MAXIMUM:	128	124	123	81	82	91	83	81	88	104	89	80	66	55	5	96	87	83	93	146	176	130	142	165	176	38586	

Figure 44. The Usage Report—Tape Mounts (continued)

9 The HOURLY MAXIMUM row shows the peak number of mounts for each hour of the reporting period.

10 This total is the maximum number of mounts for any one hour of the time period of the volume mount analyzer study. If you scan up the MAXIMUM column, you can see that the maximum (176) occurred on 6/25/91 in hour 20.

Figure 42 on page 99 shows that there was a maximum of 15 drives allocated during hour 20 on 6/25/91. Dividing 176 by 15 results in an average of 11.7 mounts per drive per hour, which is approximately one mount every 5 minutes.

The Volume Report

The Volume (VOL) Report is an actual report. The actual reports show the current tape usage statistics for your installation.

Use the keyword **REPORT(VOLUME)** or **REP(VOL)** to generate the Volume Report. You can specify the **CUTOFF** keyword to reduce the size of the Volume Report.

In this example, these keywords were specified:

```
REPORT (VOLUME) , VSORT (VOLUME)
```

These keywords specify that GFTAVMA generates a Volume Report to show volumes sorted by volume name, which is a 6-character volume serial number.

Keywords that Affect the Volume Report

The results of the GFTAVMA reports rely on the keywords you choose. Therefore, choose your keywords carefully and be aware of the defaults and effects of the keywords on these reports.

All filtering input keywords and the **SPLIT** keyword affect these reports. See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 and “SPLIT Keyword” on page 151 for the syntax and functions of these keywords.

A subset of the tailoring output keywords specified for a GFTAVMA run also affects the Volume Report. See Chapter 7, “Using GFTAVMA Keywords to Tailor Your Output,” on page 131 for more information about the syntax and functions of these tailoring output keywords:

- BLKSIZE
- CUTOFF
- DASDDEV
- FORCECAP
- RPROGRAM
- VSORT

See Appendix A, “A Quick Reference to Volume Mount Analyzer Keywords,” on page 157 and Appendix B, “Volume Mount Analyzer Keyword Syntax,” on page 159 for more information about GFTAVMA keywords.

Use of the Volume Report

The detailed Volume Report provides further information about each tape volume and its contents to supplement the data set statistics in the GFTAVMA summary report, GFTASRT2.

Usually, the GFTASRT2 summary report provides enough volume information (Figure 17 on page 54 and Figure 18 on page 56). However, if you need detailed information at the volume level, you must request the Volume Report.

The Volume Report shows your tape usage from the volume perspective. It shows the amount of data transferred to an individual volume during the period of the volume mount analyzer study.

Therefore, you can use the Volume Report to determine the efficiency of your current tape usage. If 150MB of data are transferred to a 200MB volume, your usage will be efficient. However, the typical usage is 5 or 10MB of data transferred to a 200MB volume, which is not efficient. Note that the Volume Report does not show you how to improve the inefficiency of your volume usage; it just shows whether or not your volumes are used inefficiently.

You can also use the Volume Report to see how many data sets are being placed on a volume and whether or not the volumes are full for a specific application. However, you cannot create a filter list based on the Volume Report. Volume serials are not useful tape mount management filters because data sets move from tape to tape.

The Volume Report Legend—Explanation of Symbols

Figure 45 is the header for the Volume Report. It explains some of the Volume Report column headings.

GFTASRT2 -- VOLUME REPORT -- FIELD KEY CHART

***** HEADING = 'VOLUME NAME' ***** **1**

```

+-----+
|  VOLUME  |
|  NAME    |
+-----+
  VVVVVV      (VOLUME SERIAL NUMBER) 2
FIELD-->12
3

```

FIELD	VALUES	MEANING
#1-->	>	- THIS VOLUME NO LONGER HAS ANY VALID DATA ON IT AND IT HAS BEEN RETURNED TO THE 'SCRATCH' POOL
#2-->	*	- THE DATA ON THIS DEVICE HAS EXCEEDED THE EFFECTIVE CAPACITY AND 'FORCECAP' WAS SPECIFIED 4
	X	- THE DATA ON THIS DEVICE HAS EXCEEDED THE EFFECTIVE CAPACITY AND 'NOFORCECAP' WAS SPECIFIED

***** HEADING = 'LABEL TYPE ' ***** **5**

```

+-----+
|  LABEL  |
|  TYPE   |
+-----+
  LLLLL      (VOLUME LABEL TYPE)

```

FIELD	VALUE	MEANING
LLLLL-->	SL	- THE VOLUME HAS 'STANDARD' TAPE LABELS
	NSL	- THE VOLUME HAS 'NON-STANDARD' TAPE LABELS
	AL	- THE VOLUME HAS 'ISO/ANSI' TAPE LABELS
	NL	- THE VOLUME DOESN'T HAVE ANY TAPE LABELS
	UNKN	- THE VOLUME LABEL TYPE CAN'T BE DETERMINED BY THE GFTAVMA PROGRAM

***** HEADING = 'SZ-MB XFERD' ***** **6**

```

+-----+
|  SZ-MB  |
|  XFERD  |
+-----+
  MMMMMMMM      (# OF MB ON EACH VOLUME THE LAST TIME IT WAS WRITTEN TO)
FIELD-->1

```

FIELD	VALUES	MEANING
#1-->	*	- THE 'SZ-MB' AMOUNT WAS FORCED TO EQUAL THE TOTAL NATIVE CAPACITY OF THE DEVICE BECAUSE OF THE 'FORCECAP' OPTION
	X	- ALTHOUGH THE 'SZ-MB' OF DATA ON THIS VOLUME EXCEEDS THE CAPACITY, THE VALUE WAS NOT MODIFIED DUE TO 'NOFORCECAP'

Figure 45. The Volume Report Legend

1 This refers to the column of the report with the heading Volume Name.

2 The third through eighth positions of the Volume Name field contain the Volume Serial Number, but not other symbols.

Volume Report

3 Fields 1 and 2 of Volume Name might contain symbols as described in **4**.

4 List of possible values that could appear in the indicated field.

5 This refers to the column of the report with the heading Label Type. This section of the heading explains the meanings of the five possible values that can appear in the Label Type column. These values are standard label (SL), non-standard label (NSL), ISO/ANSI label (AL), no label (NL), and unknown label type (UNKN). For more information on label types, see *z/OS DFSMS: Using Magnetic Tapes*.

6 This refers to the column of the report with the heading SZ-MB XFERD, which shows the size in megabytes of the data transferred to each volume.

Example of the Volume Report

The example report included here is only *part* of the complete Volume Report that GFTAVMA generated. Since this is only a partial Volume Report, the volume numbers start at 240.

Figure 46 on page 105 shows the Volume Report sorted according to volume using the **VSORT(VOLUME)** keyword.

GFTASRT2 -- VOLUME REPORT

7		8			9				10				
MRV	VOL#	VOLUME NAME	UNIT TYPE	LABEL TYPE	CUML MOUNTS	MOUNTS	CUML-GB XFERD	SZ-MB XFERD	3390-3 CUML-VOLS	3390-3 TRACKS	#DSNS	# INP OPENS	# OUT OPENS
120.0	240	*057686	3420	SL	2391	1	3.8	120*	2.2	3666	1	1	0
70.9	241	059081	3420	SL	2392	1	3.9	71	2.3	2217	1	1	0
4.0	242	059083	3420	SL	2393	1	3.9	4	2.3	125	1	1	0
6.1	243	059084	3420	SL	2394	1	3.9	6	2.3	190	1	1	0
2.6	244	059086	3420	SL	2395	1	3.9	3	2.3	80	1	1	0
120.0	245	*059113	3420	SL	2396	1	4.0	120*	2.3	3666	1	1	0
119.9	246	059128	3420	SL	2397	1	4.2	120	2.4	3667	1	1	0
4.0	247	059131	3420	SL	2398	1	4.2	4	2.4	123	1	1	0
2.1	248	059181	3420	SL	2400	2	4.2	4	2.4	90	1	1	1
0.0	249	063442	3420	SL	2401	1	4.2	< 1	2.4	1	1	0	1
0.1	250	071425	348X	UNKN	2402	1	4.2	< 1	2.4	3	1	1	0
2.9	251	111111	3420	SL	2403	1	4.2	3	2.4	69	1	0	1
0.0	252	154634	348X	SL	2406	3	4.2	< 1	2.4	1	1	17	0
0.1	253	403426	348X	UNKN	2407	1	4.2	< 1	2.4	4	1	1	0
0.0	254	500267	348X	SL	2408	1	4.2	< 1	2.4	22	1	1	0
19.5	255	700001	348X	SL	2409	1	4.2	19	2.4	412	1	0	1
184.5	256	700002	348X	SL	2412	3	4.7	553	2.8	16894	1	2	3
295.1	257	700004	348X	SL	2413	1	5.0	295	3.0	9221	25	0	25
24.4	258	700005	348X	SL	2414	1	5.1	24	3.0	745	1	0	1
75.9	259	700006	348X	SL	2417	3	5.3	228	3.1	6946	1	0	3
154.0	260	700009	348X	SL	2418	1	5.4	154	3.2	4715	1	1	1
17.6	261	700011	348X	SL	2420	2	5.5	35	3.2	730	1	0	5
10.9	262	700015	348X	SL	2424	4	5.5	44	3.2	1335	1	1	7
0.0	263	700016	348X	SL	2426	2	5.5	< 1	3.2	1	1	0	6
102.6	264	700019	348X	SL	2429	3	5.8	308	3.4	9395	1	0	12
2.2	265	700022	348X	SL	2433	4	5.8	9	3.4	203	1	2	2
0.1	266	700023	348X	SL	2436	3	5.8	< 1	3.4	8	10	0	13
0.0	267	700025	348X	SL	2438	2	5.8	0	3.4	0	1	1	3
73.7	268	700026	348X	SL	2440	2	6.0	147	3.5	4501	1	0	6
57.2	269	700027	348X	SL	2442	2	6.1	114	3.6	3495	1	0	2
9.5	270	700028	348X	SL	2444	2	6.1	19	3.6	595	1	0	2
250.5	271	700030	348X	SL	2446	2	6.6	501	3.9	15295	1	0	2
75.8	272	700031	348X	SL	2449	3	6.9	227	4.0	6937	1	0	9
0.7	273	700032	348X	SL	2451	2	6.9	1	4.0	30	1	0	2
9.3	274	700034	348X	SL	2453	2	6.9	19	4.1	386	1	0	6
89.7	275	700035	348X	SL	2456	3	7.1	269	4.2	8214	1	0	72
1.7	276	700036	348X	SL	2460	4	7.1	7	4.2	145	11	0	15
0.0	277	700037	348X	SL	2462	2	7.1	< 1	4.2	1	1	0	2
0.5	278	700039	348X	SL	2465	3	7.1	2	4.2	46	23	0	26
0.0	279	700040	348X	SL	2472	7	7.1	< 1	4.2	5	2	1	7
1.2	280	700041	348X	SL	2477	5	7.2	6	4.2	177	6	2	9
0.0	281	700042	348X	SL	2481	4	7.2	< 1	4.2	2	1	1	16
0.3	282	700045	348X	SL	2482	1	7.2	< 1	4.2	115	1	1	0
0.2	283	700046	348X	SL	2484	2	7.2	< 1	4.2	15	1	0	2
6.0	284	700047	348X	SL	2490	6	7.2	36	4.3	1196	1	3	4
37.0	285	700048	348X	SL	2493	3	7.3	111	4.3	3385	1	0	4
0.1	286	700049	348X	SL	2495	2	7.3	< 1	4.3	4	1	0	3
56.9	287	700050	348X	SL	2497	2	7.4	114	4.4	3471	1	0	2
2.7	288	700051	348X	SL	2501	4	7.4	11	4.4	333	1	3	1
0.0	289	700053	348X	SL	2502	1	7.4	< 1	4.4	1	1	0	1
111.0	290	700055	348X	SL	2505	3	7.8	333	4.6	10164	1	0	3

Figure 46. Partial Volume Report

7 Mount ratio value (MRV) is the ratio of the amount of data on the volume to the number of mounts attributed to that volume (MRV = number of megabytes on volume / number of mounts). A volume that has a low MRV is mounted more frequently and contains a smaller amount of data than other volumes.

8 The unit type comes from the UCB. 348X is either a 3480 with IDRC or a 3490. The label type is explained in the Volume Report Legend.

9 Number of times a volume was mounted. For example, data set A resides on tape volume AAAA at the beginning of the month. Data set B resides on AAAA in the middle of the month. Data set C resides on AAAA at the end of the month. At the end of the month, the Volume Report shows that the unique volser AAAA has mounts for data sets A, B, and C. Therefore, the totals show one volume with many mounts.

10 The target DASD device used in these calculations is controlled by the DASDDEV keyword, which defaulted to 3390-3 in this example.

11 Total size of all data sets on the volume. The size of each data set is the size when the data set was last written to the volume. Some of the values for the size of megabytes transferred are less than one (< 1).

The asterisk (*) in the SZ-MB XFERD column indicates that the number of megabytes transferred was forced to equal the total native capacity of the device because of **FORCECAP**. See “FORCECAP/NOFORCECAP Keywords” on page 140 for more details.

12 Number of DASD tracks this amount of data represents. This calculation takes into account the LRECL, RECFM, and block size of the data, and the **DASDDEV** specification.

13 Number of data sets on the volume. As this example shows, most volumes contain only one data set.

If you determine that you are using many of your volumes inefficiently, as this installation is, you can use tape mount management to stack the small data sets and fill the volumes.

Interpreting the Simulation Reports

The simulation reports show how you can change the tape usage at your installation using tape mount management. Using these reports you can simulate what your tape usage would be if you used DFSMS to manage your storage and if you used various tape devices and compaction.

There are three simulation reports:

- The Estimate Report (see “The Estimate Report”)
- The Improved Data Recording Capability (IDRC) Report (see “The Improved Data Recording Capability Report” on page 120)
- The Management Class (MGMTCLAS) Report (see “The Management Class Report” on page 125)

The Estimate Report

The Estimate Report is a simulation report. The simulation reports show how you might change the tape usage at your installation using tape mount management.

Use the keyword **REPORT(ESTIMATE)** or **REP(EST)** to generate the Estimate Report.

If you specify:

```
REPORT (EST) TAPEDEV (3490E) TAPELEN (2)
```

GFTAVMA will generate an Estimate Report that models a tape mount management environment using the 3490E tape devices and the Enhanced Capacity Cartridge System Tape.

The defaults for the keywords that were not specified are:

- **BLKSIZE** defaulted to **32 760**
- **COMPACTION** defaulted to **50%**
- **DASDDEV** defaulted to **3390-3**
- **IDRCFACTOR** defaulted to **4**
- **LARGE** defaulted to **600MB**
- **LOAGE** defaulted to **1**

- **L1AGE** defaulted to 1
- **FSPACE** defaulted to 5%
- **L1FSPACE** defaulted to 5%

See Chapter 7, “Using GFTAVMA Keywords to Tailor Your Output,” on page 131 for keyword syntax.

Note: Do not specify **GDG(GROUP)** when you request the Estimate Report. (See “GDG Keyword” on page 142 more information). Using the **DATE** and **TIME** keywords when you request the Estimate Report might impact the accuracy of the report.

Keywords that Affect the Estimate Report

The results of the GFTAVMA reports rely on the keywords you choose. Therefore, choose your keywords carefully and be aware of the defaults and effects of the keywords on these reports.

All filtering input keywords and the **SPLIT** keyword affect these reports. See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 and “SPLIT Keyword” on page 151 for the syntax and functions of these keywords.

A subset of the tailoring output keywords specified for a GFTAVMA run also affects the Estimate Report. See Chapter 7, “Using GFTAVMA Keywords to Tailor Your Output,” on page 131 for more information about the syntax and functions of these tailoring output keywords:

- BLKSIZE
- BPROGRAM
- CLASSIFY
- COMPACTION
- DASDDEV
- FORCECAP
- FSPACE
- IDRCFACTOR
- LARGE
- LEVEL0AGE
- LEVEL1AGE
- L1FSPACE
- RPROGRAM
- TAPEDEV
- TAPELEN

See Appendix A, “A Quick Reference to Volume Mount Analyzer Keywords,” on page 157 and Appendix B, “Volume Mount Analyzer Keyword Syntax,” on page 159 for more information about GFTAVMA keywords.

Description

The Estimate Report consists of three breakdown reports, a legend that explains the report’s assumptions, and a savings/cost summary that is the bottom line for the savings and costs of implementing tape mount management.

Use of the Estimate Report

The Estimate Report simulates a tape mount management environment to determine the costs and savings associated with implementing the tape mount management methodology. Use the Estimate Report to determine whether the savings of tape mount management outweigh the DASD costs.

Estimate Report

You can run the Estimate Report for any data subset because the report is generated *after* filtering.

Run the Estimate Report and the Management Class Report together to help you determine the best values for **LEVEL0AGE** and **LEVEL1AGE**. See “The Management Class Report” on page 125 for more information.

Refer to *z/OS DFSMS: Implementing System-Managed Storage* and *MVS/ESA SML: Managing Data* for additional information on the Estimate Report and the Management Class Report.

The Estimate Report—Statistical Mount Savings Breakdown

The Estimate Report includes three breakdown reports to supplement the Savings/Cost Summary (Figure 51 on page 118). The Statistical Mount Savings Breakdown provides the details for the Statistical Mount Savings reported in the Savings/Cost Summary. The total number of current statistical mounts is arranged by tape mount management data set category, and savings are attributed to the particular management class that would have been assigned based on the specified values for **L0AGE** and **L1AGE**.

The first breakdown report is the Statistical Mount Savings Breakdown Report shown in Figure 47.

GFTASRT3 -- ESTIMATE REPORT -- 3 DAY SAMPLE

STATISTICAL MOUNT SAVINGS BREAKDOWN										
DATA SET CATEGORY	CURR STAT MOUNTS	TARG STAT MOUNTS	DFSMS - PRIMARY		HSM - LEVEL 1		HSM - LEVEL 2		TAPE - DIRECT	
			CURR MNT#	TARG MNT#	CURR MNT#	TARG MNT#	URR MNT#	TARG MNT#	CURR MNT#	TARG MNT#
TEMPORARY	242.0		242.0							
ACTIVE	2088.5	9.9	1483.1		0.0		605.4	9.9		
ACTV GDG	6443.4	132.0	1930.6		0.0		4512.8	132.0		
BACKUP	25476.8	1131.5	7916.6				17560.2	1131.5		
LARGE	4050.0	2380.2							4050.0	2380.2
HSM SFF	0.0	0.0							0.0	0.0
HSM GEN	0.0	0.0							0.0	0.0
MGMTCLAS		3384.0						3384.0		
T O T A L	38300.8	7037.6	11572.3		0.0		22678.4	4657.4	4050.0	2380.2

Figure 47. Estimate Report—Statistical Mount Savings Breakdown

1 In this example, the current statistical mounts value (CURR STAT MOUNTS) for active data sets is 2088.5. The target statistical mounts (TARG STAT MOUNTS) is set to the HSM - LEVEL 2 TARG MNT# for the ACTIVE, ACTV GDG, and BACKUP categories, and the numbers in these columns are the same.

The estimated target statistical mounts value will be 9.9 if you use DFSMS to manage tape data sets. This reduction results from managing tape data sets in the storage hierarchy using 3490E tape hardware. The reduction from 2088.5 to 9.9 is significant because the data sets representing 1483.1 of the mounts are residing on primary DASD, and the data sets representing the remaining 605.4 mounts are on migration level 2 tape. Because of 3490E and, optionally, DFSMSHsm compaction, these mounts are reduced to 9.9 target mounts.

See “Statistical Mounts” on page 30 for more information.

2 For ACTV GDG data sets, the values in the HSM - LEVEL 2 TARG MNT# column in Figure 47 on page 108 are equal to the values in the HSM - LEVEL 2 TARG VOL# column in Figure 48 on page 111 because when the ACTV GDG volumes are mounted, all the data sets are written to these volumes during one output operation. Thus, full volumes are written.

3 GFTAVMA models large data sets (greater than the **LARGE** keyword default of 600MB) as going directly to migration level 2 tape. Therefore, DFSMS would not manage them, and they would never reside in the storage hierarchy. The mount reduction from 4050.0 to 2380.2 is based on modelling to the 3490E with the Enhanced Capacity Cartridge System Tape (**TAPEDEV** of 3490E and **TAPELEN** of 2), which has a much larger capacity.

4 The target mounts incurred for migration level 2 recalls are reported in the MGMTCLAS row.

5 In this example, the volume mount analyzer simulates an environment where DFSMSHsm bypasses migration level 1 processing. Therefore, the columns for migration level 1 contain zeros. The data sets are modelled as going directly from primary DASD to migration level 2 tape, bypassing migration level 1 DASD because the **LEVEL0AGE** and the **LEVEL1AGE** keywords defaulted to one.

6 The target HSM migration level 2 mounts for ACTIVE, ACTV GDG, and BACKUP represent the mounts needed to create migration level 2 tape volumes.

Understanding the Columns of the Report

The columns in the Statistical Mount Savings Breakdown are:

DATA SET CATEGORY:

The tape mount management data set categories. See “Categories of Tape Mount Management Data Sets” on page 38 for a description of these categories.

CURR STAT MOUNTS:

The current number of statistical mounts for each tape mount management data set category. See “Statistical Mounts” on page 30.

TARG STAT MOUNTS:

The number of statistical mounts that would still occur for each tape mount management category if tape mount management was implemented according to the various parameters specified in the GFTAVMA run.

DFSMS - PRIMARY CURR MNT#:

The number of mounts that occurred because the corresponding data sets are residing on DFSMS *primary storage*.

DFSMS - PRIMARY TARG MNT#:

This field is always blank because these values are always zero. The PRIMARY TARG MNT# value is reduced to zero because the data sets that caused these mounts would still be residing on primary DASD in the simulated environment.

HSM - LEVEL 1 CURR MNT#:

The number of mounts that occurred because the corresponding data sets are residing on DFSMSHsm migration level 1 storage.

HSM - LEVEL 1 TARG MNT#:

This field is always blank because no physical tape mounts are required to recall a data set that is compacted on migration level 1 DASD.

Estimate Report

HSM - LEVEL 2 CURR MNT#:

The number of mounts that occurred because the corresponding data sets are residing on DFSMSHsm migration level 2 storage.

HSM - LEVEL 2 TARG MNT#:

Number of mounts that would be incurred by DFSMSHsm migration level 2 in the target environment.

TAPE DIRECT CURR MNT#:

Number of current mounts that fall into the Tape Direct Category (for example, data sets larger than the **LARGE** keyword value).

TAPE DIRECT TARG MNT#:

Number of mounts for the **LARGE** data sets in the target environment using the type of tape hardware specified by the **TAPEDEV** and **TAPELEN** keywords.

Understanding the Rows of the Report

The first column of the breakdown reports is **DATA SET CATEGORY**. The rows are **TEMPORARY**, **ACTIVE**, **ACTV GDG**, **BACKUP**, **LARGE**, **HSM SFF**, **HSM GEN**, and **MGMTCLAS**. These categories are related to the tape mount management data categories. See “Categories of Tape Mount Management Data Sets” on page 38 and “Data Set Usage Types” on page 36 for descriptions of these categories.

The tape mount management category of **Active** is divided into **ACTIVE** and **ACTV GDG** for the Estimate Report.

The tape mount management category of **HSM** is divided into **HSM SFF** (DFSMSHsm single file format data sets) and **HSM GEN** (DFSMSHsm generated data sets, such as **VTOCDUMPs**).

MGMTCLAS is the row that shows *recall mounts*.

The Estimate Report—Statistical Volume Savings Breakdown

The Statistical Volume Savings Breakdown provides the details for the Statistical Volume Savings reported in the Savings/Cost Summary (Figure 51 on page 118). The total number of current statistical volumes is arranged by tape mount management data set category, and savings are attributed to the particular management class that would have been assigned based on the specified values for **L0AGE** and **L1AGE**.

Figure 48 on page 111 is the second breakdown report. This report is identical to the Statistical Mounts Savings Breakdown report except that the values are **STATISTICAL VOLUMES** instead of **STATISTICAL MOUNTS**.

STATISTICAL VOLUME SAVINGS BREAKDOWN										
DATA SET CATEGORY	CURR STAT VOLUMES	TARG STAT VOLUMES	DFSMS - PRIMARY CURR VOL#	TARG VOL#	HSM - LEVEL 1 CURR VOL#	TARG VOL#	HSM - LEVEL 2 CURR VOL#	TARG VOL#	TAPE - DIRECT CURR VOL#	TARG VOL#
TEMPORARY	48.2		48.2							
ACTIVE	86.8	9.9	18.7		0.0		68.1	9.9		
ACTV GDG	635.4	132.0	90.4		0.0		545.0	132.0		
BACKUP	4987.3	1131.5	760.7				4226.6	1131.5		
LARGE	557.8	323.8						557.8	323.8	
HSM SFF	0.0	0.0						0.0	0.0	
HSM GEN	0.0	0.0						0.0	0.0	
MGMTCLAS		4.2						4.2		
T O T A L	6315.5	1601.4	918.0		0.0		4839.7	1277.6	557.8	323.8

Figure 48. Estimate Report—Statistical Volume Savings Breakdown

7 In this example, the current statistical volumes value (CURR STAT VOLUMES) for active data sets is 86.8. The target statistical volumes (TARG STAT VOLUMES) is set to the HSM - LEVEL 2 TARG VOL# for the ACTIVE, ACTV GDG, and BACKUP categories, so these numbers are the same.

The estimated target statistical volumes value if you used DFSMS to manage tape data sets would be 9.9.

This reduction from 86.8 to 9.9 is significant because 18.7 of the 86.8 volumes of active data sets are still residing on primary DASD. The rest of the volumes (68.1) have been reduced to 9.9 volumes because of managing tape data sets in the storage hierarchy and using 3490E tape hardware and Enhanced Capacity Cartridge System Tapes. See “Statistical Volumes” on page 31 for more details.

8 For ACTV GDG data sets, the values in the HSM - LEVEL 2 TARG VOL# column should be equal to the HSM - LEVEL 2 TARG MNT# because when the ACTV GDG volumes are mounted, all the data sets are written to these volumes during one output operation. Thus, full volumes are written.

9 GFTAVMA models large data sets (greater than the **LARGE** default of 600MB) as going directly to tape, so DFSMS does not manage them in the storage hierarchy. The volume reduction from 557.8 to 323.8 is based on modelling to the 3490E with the Enhanced Capacity Cartridge System Tape (TAPEDEV of 3490E and TAPELEN of 2), which has a much larger capacity.

10 In this example, the volume mount analyzer simulates an environment where DFSMSHsm bypasses migration level 1 processing. Therefore, the columns for migration level 1 contain zeros. GFTAVMA models the data sets as going directly from primary DASD to migration level 2 tape, bypassing migration level 1 DASD because the **LEVEL0AGE** and the **LEVEL1AGE** keywords defaulted to one.

Understanding the Columns of the Report

The columns in the Statistical Volume Savings Breakdown are:

DATA SET CATEGORY:

The tape mount management data set categories. See “Categories of Tape Mount Management Data Sets” on page 38 for a description of these categories.

CURR STAT VOLUMES:

The current number of statistical volumes for each tape mount management data set category. See “Statistical Volumes” on page 31 for more details.

TARG STAT VOLUMES:

The number of statistical volumes that would still occur in each tape mount management data set category if tape mount management was implemented according to the various parameters specified in the GFTAVMA run.

DFSMS - PRIMARY CURR VOL#:

The number of volumes used because the corresponding data sets (such as temporary data sets) are residing on DFSMS primary storage.

DFSMS - PRIMARY TARG VOL#:

This field is always blank because these values are always zero. The PRIMARY TARG VOL# is reduced to zero because the data sets representing these volumes are still residing on primary DASD for the simulation.

HSM - LEVEL 1 CURR VOL#:

The number of volumes used because the corresponding data sets are residing on DFSMSHsm migration level 1 storage.

HSM - LEVEL 1 TARG VOL#:

This field is always blank.

HSM - LEVEL 2 CURR VOL#:

The number of volumes used because the corresponding data sets are residing on DFSMSHsm migration level 2 storage.

HSM - LEVEL 2 TARG VOL#:

Number of volumes that would be used by DFSMSHsm migration level 2 in the target environment.

Note: The difference between the current and target values is the result of stacking data sets (because DFSMSHsm fills the tapes) and the new tape technology for which you are estimating if it is different from your current technology.

TAPE DIRECT CURR VOL#:

Number of current volumes that fall into the Tape Direct Category (for example, data sets larger than the **LARGE** keyword value).

TAPE DIRECT TARG VOL#:

Number of volumes for the **LARGE** data sets in the target environment using the type of tape hardware specified by the **TAPEDEV** and **TAPELEN** keywords.

Understanding the Rows of the Report

The first column of the breakdown reports is DATA SET CATEGORY. The rows are TEMPORARY, ACTIVE, ACTV GDG, BACKUP, LARGE, HSM SFF, HSM GEN, and MGMTCLAS. These categories are related to the tape mount management data categories. See “Categories of Tape Mount Management Data Sets” on page 38 and “Data Set Usage Types” on page 36 for descriptions of these categories.

The tape mount management category of Active is divided into ACTIVE and ACTV GDG for the Estimate Report.

The tape mount management category of HSM is divided into HSM SFF (DFSMSHsm single file format data sets) and HSM GEN (DFSMSHsm generated data sets, such as VTOCDUMPs).

MGMTCLAS is the row that shows recall mounts.

The Estimate Report—DASD Buffer Cost Breakdown

The tape mount management DASD Buffer Cost Breakdown provides the details for the DASD buffer costs in the Savings/Cost Summary (Figure 51 on page 118). This breakdown report shows the distribution of data sets by tape mount management data set category and the way they will reside in the simulated target environment.

Figure 49 is the third breakdown report. All the values in this report represent *weighted averages* based on the number of days the data sets resided on the various levels of storage. For example, a data set with 500 tracks has a size of 15MB on migration level 2. If it resides on migration level 2 for two thirds of the sample, the DASD Buffer Cost Breakdown shows it having a value of 10MB.

GFTASRT3 -- ESTIMATE REPORT -- 33 DAY SAMPLE

11

DASD BUFFER COST BREAKDOWN											
DATA SET CATEGORY	DSN#	ORIGINAL SIZE-MB	ORIGINAL SIZE-TK	DFSMS - PRIMARY DSN#	DFSMS - PRIMARY SIZE-TK	HSM - LEVEL 1 DSN#	HSM - LEVEL 1 SIZE-TK	HSM - LEVEL 2 DSN#	HSM - LEVEL 2 SIZE-MB	TAPE - DIRECT DSN#	TAPE - DIRECT SIZE-MB
TEMPORARY	141	3039	83217	6	659						
ACTIVE	200	6779	196782	29	33999	0	0	171	3942		
ACTV GDG	1685	71285	1977258	104	120649	0	0	1581	52769		
BACKUP	38794	761995	20359815	2030	1095199			36764	451926		
LARGE	418	466686	14200124							418	203378
HSM SFF	0	0	0							0	0
HSM GEN	0	0	0							0	0
MGMTCLAS					103622		0		1689		
T O T A L	41238	1309784	36817196	2169	1354128	0	0	38516	506948	418	203378

15

Figure 49. Estimate Report—DASD Buffer Cost Breakdown

11 In some cases, the HSM LEVEL 1 SIZE-TK column might contain values even if the HSM LEVEL 1 DSN# values are zero. If data has resided on migration level 1 at some time during this sample, there could be zeros in the DSN# column and values for the number of tracks used in the SIZE-TK column.

12 There are 1885 active data sets in this sample, 200 non-GDG data sets and 1685 GDG data sets. Based on the keywords specified or defaulted (see “Estimating Savings and Costs” on page 117) in this GFTAVMA run, 29 of the active non-GDSs would reside on the tape mount management DASD buffer (DFSMS DSN# PRIMARY SIZE-TK) while 171 of them would have already migrated to migration level 2 tape (HSM LEVEL 2 DSN#).

13 The TAPE DIRECT column shows the LARGE data sets, which DFSMS would not intercept, but would write directly to tape. In this case, 418 data sets are greater than the LARGE keyword value, so they would be modelled as going directly to tape. The size reduction from 466686MB to 203378MB would be a result of IDRCFACTOR of 3, TAPEDEV of 3490E, and TAPELEN of 2.

14 The MGMTCLAS row shows the data set recall activity caused by your selection of LEVEL0AGE and LEVEL1AGE. You can reduce the recall mounts by changing the values for L0AGE and L1AGE. See “The Management Class Report” on page 125 for details on changing values.

15 In this example, the volume mount analyzer simulates an environment where DFSMSHsm bypasses migration level 1 processing. Therefore, the columns for migration level 1 contain zeros. The data sets go directly from primary DASD to migration level 2 tape, bypassing migration level 1 DASD because the LEVEL0AGE and LEVEL1AGE keywords defaulted to one.

If the L1AGE value is greater than L0AGE (for example, L0AGE=1 and L1AGE=3), then the data sets that are not referenced would reside on primary DASD for one day, then migrate to migration level 1 DASD and reside for an additional two days, and finally migrate to migration level 2 tape. They would go to migration level 2 tape after a total of three days since the data sets were last referenced.

Understanding the Columns of the Report

The columns in the DASD Buffer Cost Breakdown are:

DATA SET CATEGORY:

The tape mount management data set categories. See “Categories of Tape Mount Management Data Sets” on page 38 for a description of these categories.

DSN#:

A breakdown of the total number of data sets for each tape mount management data set category represented in the input SMF sample.

ORIGINAL SIZE-MB:

The amount of data transferred in megabytes for each tape mount management data set category.

ORIGINAL SIZE-TK:

If all data sets in the tape mount management categories were placed on a *DASD volume* and not migrated or deleted, this column would show the total number of DASDDEV tracks required to hold all the data sets.

DFSMS—PRIMARY:

The primary storage estimate determined by the values specified for the DASDDEV and LEVEL0AGE keywords and by the reference activity of the data sets.

DSN#:

The number of data sets by tape mount management category that would reside in primary storage.

SIZE-TK:

The number of DASDDEV tracks that would be required for primary storage. This size estimate is the weighted average size of the primary storage required for the entire reporting period represented in the SMF sample.

HSM—LEVEL 1:

The DFSMSHsm migration level 1 storage estimate determined by the values

specified for the **DASDDEV**, **COMPACTION**, and **LEVEL1AGE** keywords and by the reference activity of the data sets.

DSN#:

The number of data sets by tape mount management category that would reside on migration level 1 storage.

SIZE-TK:

The number of **DASDDEV** tracks that would be required for migration level 1 storage. This size estimate is the weighted average size of the migration level 1 storage required for the entire reporting period represented in the SMF sample.

HSM—LEVEL 2:

The DFSMSHsm migration level 2 storage estimate determined by the values specified for the **TAPEDEV**, **TAPELEN**, **IDRC** (or the current **IDRCFACTOR** for the data set), and **COMPACTION** keywords and by the reference activity of the data sets. This column represents tape storage for data sets that DFSMSHsm would migrate.

Note: Level 2 storage is not part of the tape mount management DASD buffer.

DSN#:

The number of data sets by tape mount management category that would reside on migration level 2 tape storage.

SIZE-MB:

The number of megabytes (MB) that would be required for migration level 2 tape storage. This size estimate is the weighted average size of the migration level 2 storage required for the entire period represented in the SMF sample.

TAPE DIRECT:

The tape storage estimate for the data sets that would continue to go directly to tape and not be intercepted by DFSMS. This estimate is determined by the **LARGE**, **TAPEDEV**, **TAPELEN**, and **IDRC** (or the current **IDRCFACTOR** for the data set) keywords.

Note: This storage is for tape requirements and is not part of the tape mount management DASD buffer.

DSN#:

The number of data sets that would continue to go directly to tape.

SIZE-MB:

The number of megabytes that would be required for tape storage.

Understanding the Rows of the Report

The first column of the breakdown reports is **DATA SET CATEGORY**. The rows are **TEMPORARY**, **ACTIVE**, **ACTV GDG**, **BACKUP**, **LARGE**, **HSM SFF**, **HSM GEN**, and **MGMTCLAS**. These categories are related to the tape mount management data categories. See “Categories of Tape Mount Management Data Sets” on page 38 and “Data Set Usage Types” on page 36 for descriptions of these categories.

The tape mount management category of **Active** is divided into **ACTIVE** and **ACTV GDG** for the Estimate Report.

The tape mount management category of **HSM** is divided into **HSM SFF** (DFSMSHsm single file format data sets) and **HSM GEN** (DFSMSHsm generated data sets, such as **VTOCDUMPs**).

Estimate Report

MGMTCLAS is the row that shows recall mounts.

The Estimate Report Legend—Explanation of Symbols

The Estimate Report legend shows the specified and defaulted keywords that were used to produce the Estimate Report. It also defines the tape mount management data set categories and shows how DFSMS would manage the data sets in each of those categories in the simulated environment. See “Categories of Tape Mount Management Data Sets” on page 38 for more information about the various categories of data.

Figure 50 is the legend for this example Estimate Report. It explains the assumptions GFTAVMA used when calculating the numbers for the Estimate Report. This legend shows how GFTAVMA determines which data sets would be migrated to which storage levels.

```
*****
*
* TARGET DEVICES:          FREE SPACE:
* DASDDEV = 3390-3        FSPACE = 5%
* TAPEDEV = 3490E        L1FSPACE = 5%
* TAPELEN = 2
*
* TAPE MOUNT MANAGEMENT CATEGORIES:
*
* +-----+
* | TEMPORARY | --> LEVEL 0 DASD AND EXPIRE AFTER USAGE
* +-----+
* | ACTIVE +  | --> LEVEL0AGE = 1 DAYS
* | ACTV GDG  | --> LEVEL1AGE = 1 DAYS
* +-----+
* | BACKUP    | --> LEVEL 2 AFTER 0 DAYS (IF MIGRATION RUNS ONCE EVERY 24 HOURS)
* +-----+
* | MGMTCLAS  | --> RECALL ACTIVITY FROM LEVEL0AGE AND LEVEL1AGE VALUES
* +-----+
*
* DATA SETS ==> DIRECTLY TO TAPE:
*
* +-----+
* | LARGE     | --> DATA SETS > 600 MB
* +-----+
* | HSM SFF   | --> HSM SINGLE-FILE FORMAT DATA SETS
* +-----+
* | HSM GEN   | --> HSM-GENERATED DATA SETS
* +-----+
*
*****
```

Figure 50. The Estimate Report Legend

16 The **DASDDEV** keyword specifies the DASD device that GFTAVMA should use as the model for the tape mount management DASD buffer, which in this example defaulted to the **DASDDEV** value of 3380E.

GFTAVMA modelled to the 3490E with Enhanced Capacity Cartridge System Tape because a **TAPEDEV** of 3490E and a **TAPELEN** of 2 were specified.

When GFTAVMA models the tape mount management DASD buffer, it uses a high free space threshold as DFSMSHsm would. The GFTAVMA defaults are 5% free space on the primary DASD buffer and 5% free space on the migration level 1 buffer. If you wish to change these free space percentages, use the **FSPACE** and **L1FSPACE** keywords.

17 The tape mount management data set categories section shows when DFSMS would migrate data sets from primary storage to migration level 1 to migration level 2:

- Temporary data sets reside on primary DASD for the duration of the job and then expire, so they never migrate.

- Active data sets and active GDG data sets are controlled by the **LEVEL0AGE** and **LEVEL1AGE** values of 1, so these data sets would reside on primary storage for one day before migrating to migration level 2 tape (if the data sets are not referenced during this time).
- Backup data sets migrate to migration level 2 tape after zero days if migration runs every 24 hours. Backup data sets would reside on primary storage for less than one day. They would never reside on migration level 1 regardless of the values for **L0AGE** and **L1AGE**.
- **MGMTCLAS** data sets are data sets that have recall activity. They are recalled from migration level 2 and migration level 1 back to primary storage. The recall activity is based on the last reference to the data set, which is calculated by **GFTAXTR**, and on the values specified by the **LEVEL0AGE** and **LEVEL1AGE** keywords.

You need to select management class reference ages for active data sets that would migrate the data at optimum times. In this example, **L0AGE** and **L1AGE** defaulted to 1, but these might not be the optimum values based on the data set usage at your installation. The Management Class Report can help you determine the recommended values to use for each application. See “The Management Class Report” on page 125.

For the Estimate Report, **GFTAVMA** assumes you specified the migration level 1 compaction value that you feel you can achieve using the **DFSMSHsm** level 1 compaction option. The default is 50% compaction, which means migration level 1 is reduced by 50% from primary storage. The Estimate Report example in this section shows 50% compaction, but you can use the **COMPACTION** keyword to change that percentage.

18 **GFTAVMA** models data sets larger than the **LARGE** keyword value (default 600MB) as going directly to tape. In this way, **DFSMS** does not manage them in the storage hierarchy. **GFTAVMA** also models **DFSMSHsm** **SFF** and **GEN** data sets as going directly to tape.

Estimating Savings and Costs

The Estimate Report is a bottom line assessment of the costs and savings from implementing tape mount management. You can use the report as a high level, first cut at the potential tape mount management savings if you specify no keywords other than **REPORT(ESTIMATE)**, and let **GFTAVMA** categorize your data. In contrast, you can tailor the report by using keywords to focus on a particular subset of data and by adjusting the **L0AGE** and **L1AGE** keywords to determine where and how long the data is stored in the storage hierarchy.

Note: The DASD costs shown in the Estimate Report are based on **DFSMSHsm** running space management against the tape mount management DASD buffer once a day. The DASD requirement could be significantly reduced by running interval migration against those volumes on an hourly basis. However, the volume mount analyzer cannot model this reduction directly.

See *MVS/ESA SML: Managing Storage Groups* and *z/OS DFSMS: Implementing System-Managed Storage* for details on **DFSMSHsm** migration.

Savings/Cost Summary

Figure 51 on page 118 shows the part of the Estimate Report that explains the bottom line for the details in the breakdown reports. The three parts of the summary reflect the three breakdown reports:

Estimate Report

1. Part **ONE** summarizes the Statistical Mount Savings Breakdown. It shows that if you used DFSMS and the 3490E tape device to intercept tape data sets, you could save 81.6% of your mounts.
2. Part **TWO** summarizes the Statistical Volume Savings Breakdown. It shows that if you used DFSMS and 3490E and Enhanced Capacity Cartridge System Tape to intercept tape data sets, you could save 74.6% of your volumes.
3. Part **THREE** summarizes the DASD Buffer Cost Breakdown. It shows that if you used DFSMS to intercept tape data sets, it would cost 28.4 3390-3 DASD volumes to eliminate 31263.2 tape mounts and 4714.1 tape volumes.

```

GFTASRT3 -- ESTIMATE REPORT -- 33 DAY SAMPLE
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     SAVINGS/COST SUMMARY                                     |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| THE FOLLOWING CHART ASSUMES THAT THE CURRENT DATE                                     |
| IS THE LAST DAY OF THE INPUT SAMPLE AND HSM IS                                     |
| SET UP TO PROCESS THE TMM DASD BUFFER ONLY ONCE A                                  |
| DAY -- THE ACTUAL TMM DASD BUFFER MAY BE LESS IF                                  |
| HSM PROCESSING IS PERFORMED ON AN HOURLY BASIS                                   |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| STATISTICAL MOUNT SAVINGS:                                                         |
| 1669.8 MOUNTS (LARGE) - 4.4%                                                       |
| 0.0 MOUNTS (HSM) - 0.0%                                                            |
| 32977.3 MOUNTS (DFSMS) - 86.1%                                                     |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 34647.2 MOUNTS - 90.5%                                                             |
| - 3384.0 MOUNTS (MGMTCLAS) - 8.8%                                                  |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 31263.2 MOUNTS (TOTAL) - 81.6%                                                     |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| STATISTICAL VOLUME SAVINGS:                                                         |
| 234.0 VOLUMES (LARGE) - 3.7%                                                       |
| 0.0 VOLUMES (HSM) - 0.0%                                                          |
| 4484.3 VOLUMES (DFSMS) - 71.0%                                                    |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 4718.4 VOLUMES - 74.7%                                                             |
| - 4.2 VOLUMES (MGMTCLAS) - 0.1%                                                  |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 4714.1 VOLUMES (TOTAL) - 74.6%                                                     |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| DASD BUFFER COST:                                                                  |
| 25.0 3390-3 VOLUMES (PRIMARY)                                                       |
| 2.1 3390-3 VOLUMES (MGMTCLAS)                                                       |
| 1.3 VOLUMES (+ 'FSPACE'====> 5%)                                                  |
| 0.0 3390-3 VOLUMES (HSM LEVEL 1)                                                  |
| 0.0 3390-3 VOLUMES (MGMTCLAS)                                                       |
| 0.0 VOLUMES (+ 'LIFSPACE'==> 5%)                                                  |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 28.4 TOTAL VOLUMES                                                                  |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

19

20

21

22

23

24

25

26

27

28

ONE

TWO

THREE

GFTAVMA PROGRAM COMPLETED AT 09:28:31 A.M. ON 06/15/2004 -- RETURN CODE = 0

Figure 51. The Estimate Report—Savings/Cost Summary

19 Data sets larger than the **LARGE** keyword value are modelled as going directly to the tape device specified by the **TAPEDEV** and **TAPELEN** keywords. DFSMS would not manage these data sets, so the 4.4% savings result from modelling to 3490E with the Enhanced Capacity Cartridge System Tape.

The projected savings shown in the Savings/Cost Summary is the difference between current statistical mounts or volumes and the target statistical mounts or volumes shown in the Estimate Report breakdown reports. For example, Figure 47 on page 108 shows the current statistical mounts and 2380.2 target statistical mounts for the **LARGE** category; the difference is a savings of 1669.8.

20 HSM mount savings result from using different tape hardware technology. DFSMSHsm mounts would be saved if you implement IDRC, 3490E, or double capacity cartridges, as specified with the **IDRCFACTOR**, **TAPEDEV**, and **TAPELEN** keywords.

21 The volume mount analyzer simulates the effects of DFSMS intercepting data sets that do not fall into the Large or HSM categories and allocating them to the tape mount management DASD buffer.

The DFSMS mounts in this part of the savings/cost summary represent the temporary, active, active GDG, and backup data sets. GFTAVMA models the migration of these data sets to migration level 1 and migration level 2 according to the specification of management class type keywords, such as **L0AGE** and **L1AGE**.

DFSMS mounts in the Savings/Cost Summary are from the TOTAL mounts row in Figure 47 on page 108. The difference between 38300.8 current mounts and 7037.6 target mounts is the savings of 31263.2.

22 GFTAVMA subtracts mounts specified as **MGMTCLAS** from the total mounts saved because these mounts represent DFSMSHsm recalls that would occur as a result of the current usage pattern. An acceptable percentage of recalls is 5%. Therefore, the 8.8% shown here indicates that you need to select different values for **L0AGE** and **L1AGE** to reduce the number of recall mounts. See “The Management Class Report” on page 125 for more information.

23 To implement tape mount management in this example, you would need 25 primary storage DASD volumes. GFTAVMA calculates the number of volumes needed by dividing the number of primary size tracks (**PRIMARY SIZE-TRK**) shown in Figure 49 on page 113 by the number of tracks on the target volumes, which in this example are 3390-3 volumes.

24 You would need 2.1 DASD volumes for recalled data sets (**MGMTCLAS**).

25 You would need 1.3 volumes for the 5% free space. Run the Maximum Gigabyte Report (see “The Maximum Gigabyte Report” on page 88) to verify that the 5% free space default is adequate for your needs.

You would need a total of 28.4 (which would be rounded up to 29) 3390-3 DASD volumes for the primary DASD buffer.

26 You would need 0 volumes for HSM migration level 1.

27 You would need 0 volumes for migration level 1 free space. You do not need any volumes for migration level 1 because, in this example, the volume mount analyzer simulates an environment where DFSMSHsm bypasses migration level 1 processing.

The data sets go directly from primary DASD to migration level 2 tape, bypassing migration level 1 DASD because the **LEVEL0AGE** and **LEVEL1AGE** keywords defaulted to one.

28 You would need a total of 28.4 (which would be rounded up to 29) 3390-3 DASD volumes for the tape mount management DASD buffer. These volumes are for primary storage space. You do not need migration level 1 space because the **LEVEL1AGE** keyword caused GFTAVMA to bypass level 1 and migrate the data directly to level 2 tape.

Estimate Report

The storage needed for the tape mount management DASD buffer and the new tape technology generate the cost for reducing volumes and mounts.

The tape mount management DASD buffer calculations in the Estimate Report assume that the data sets are migrated from the tape mount management DASD buffer by DFSMSHsm Daily Space Management. Thus, the report reflects the amount of DASD required for a minimum of 24 hours DASD residency.

Since most of the data is likely to be eligible for immediate migration, a more efficient implementation would be achieved by using DFSMSHsm interval migration processing to clear the tape mount management DASD buffer volumes on an hourly basis. The GBMAX Report helps you determine how much new data would be allocated hourly so that you can calculate the size of the DASD buffer needed for hourly migration.

See *MVS/ESA SML: Managing Storage Groups* and *z/OS DFSMS: Implementing System-Managed Storage* for details on planning your DASD buffer.

The Improved Data Recording Capability Report

The Improved Data Recording Capability (IDRC) Report is a simulation report. This simulation report shows how you can change the tape usage at your installation using IBM hardware technology.

Use the keyword **REPORT(IDRC)** or **REP(IDRC)** to generate the IDRC Report.

For this example, these keywords were specified:

```
REPORT(IDRC) LARGE (1500) TAPEDEV (3490E) IDRCFACTOR (3) TAPELEN (2)
```

These keywords specify that GFTAVMA generate an IDRC Report that shows data sets that are larger than 1500MB. The **LARGE** keyword was set to 1500 in this example to limit the report to a size that would be a good example report. Normally, you would select a **LARGE** keyword value that is the size of the tape cartridges at your installation.

Start with the **LARGE** default of 600 and increase the value if your sample is too big to be useful. For the IDRC calculations, the **TAPEDEV**, **IDRCFACTOR**, and **TAPELEN** keywords direct GFTAVMA to model to a 3490E with an IDRC compaction ratio of 3 and a tape length of 2. If you specify **LARGE (0)**, GFTAVMA evaluates all data sets as candidates for IDRC compaction.

The default for the keyword that was not specified was:

- **BLKSIZE** defaulted to **32 760**

Note: Do not specify **GDG(GROUP)** for the IDRC Report. See “GDG Keyword” on page 142 for more information.

Keywords that Affect the IDRC Report

The results of the GFTAVMA reports rely on the keywords you choose. Therefore, choose your keywords carefully and be aware of the defaults and effects of the keywords on these reports.

All filtering input keywords and the **SPLIT** keyword affect these reports. See Chapter 5, “Using GFTAVMA Keywords to Filter Your Input,” on page 63 and “SPLIT Keyword” on page 151 for the syntax and functions of these keywords.

A subset of the tailoring output keywords specified for a GFTAVMA run also affects the IDRC Report. See Chapter 7, “Using GFTAVMA Keywords to Tailor Your Output,” on page 131 for the syntax and rules for these tailoring output keywords:

- BLKSIZE
- FORCECAP
- IDRCFACTOR
- LARGE
- RPROGRAM
- TAPEDEV
- TAPELEN

See Appendix A, “A Quick Reference to Volume Mount Analyzer Keywords,” on page 157 and Appendix B, “Volume Mount Analyzer Keyword Syntax,” on page 159 for more information about GFTAVMA keywords.

Use of the IDRC Report

The IDRC Report shows the potential effects of using IDRC compaction, 3490E enhanced recording format, and the Enhanced Capacity Cartridge System Tape with large data sets. If you use IDRC compaction on a data set that resides on four volumes, you can make it fit on one volume.

To understand the numbers generated for this report, it is important to note the following:

- The IDRC Report shows the data set level detail for the large category savings reported in the Estimate Report. To determine the summary of total savings, request only the Estimate Report.

This report does not process from the volume perspective. For example, it does not show which tapes are suitable for IDRC/3490E.

- The IDRC Report only includes data sets larger than the value specified with the **LARGE** keyword. It shows the volumes and mounts associated with each of these data sets. In the tape mount management methodology, these data sets are large enough to make good use of a cartridge and would go directly to tape.

GFTAVMA assumes that DFSMS intercepts data sets smaller than the **LARGE** value (default 600MB) and re-routes them to the tape mount management DASD buffer storage group. The effects of IDRC, 3490E, and double capacity cartridges are modelled for these smaller data sets when they are eligible to be migrated to tape with DFSMSHsm, based on the values specified or defaulted for the **IDRC**, **TAPEDEV** and **TAPELEN** keywords.

- Because GFTAVMA analyzes tape processing over time, it is possible to have data sets included in the sample that no longer exist because the volumes associated with these data sets have been reused for other data. These data sets are still accounted for in the IDRC Report because, if you had used IDRC compaction at the time the data sets were written, you could have saved volumes and mounts.
- If the savings from using IDRC and 3490E are less than anticipated, it might be because the majority of your tape data sets are single volume and do not show volume or mount savings that result from using these techniques directly. You can enhance the benefits dramatically by using the principles of tape mount management to ensure that all cartridges are automatically filled with data and thus can benefit from compaction and 36-track recording.

Calculating Target Volumes and Mounts

For its simulation, GFTAVMA calculates the number of target volumes using the size of the current version of the data set. This is the size, in megabytes, that was specified in the MEGABYTES TRANSFERRED column of the IDRC Report (see Figure 53 on page 124).

Based on the **TAPEDEV**, **TAPELEN** and **IDRCFACTOR** specified, GFTAVMA calculates the effective capacity of the target volume and the number of volumes required to hold the data set. GFTAVMA calculates the projected number of target mounts using the number of target volumes, multiplied by a mount factor, which simply represents the number of times the set of volumes was used (current mounts/current volumes).

If data sets are single file, single volume or single file, multivolume (SNGL), GFTAVMA cannot compact them to less than one volume. GFTAVMA then rounds the projected number of volumes up to the next whole volume size. For example, 2½ volumes are rounded up to 3.

If data sets are multifile, single volume (MULT), GFTAVMA calculates the compaction, but does not reduce the projected volume or mount counts.

However, if data sets are multifile, multivolume (MULT), GFTAVMA does reduce their projected volume count to less than one volume because GFTAVMA assumes that the user wants to simulate the effects of IDRC on stacked data sets. Therefore, GFTAVMA reduces those data sets even though their original size was less than or equal to one statistical volume.

See “Single/Multi Volume/File Data Sets” on page 29 for more information about single/multi volume/file.

The simulated reduction of the MULT data sets is based on your adjusting the JCL to accommodate the additional capacity of the larger capacity tape cartridges.

If the data set was previously compacted with IDRC, GFTAVMA uses its actual compaction factor, and the **IDRCFACTOR** keyword does not affect the compaction. GFTAVMA only uses the **IDRCFACTOR** value if the data set was not previously compacted.

The IDRC Report Legend—Explanation of Symbols

Figure 52 on page 123 shows the report legend, which explains some of the column headings of the IDRC Report.

GFTASRT3 -- IDRC REPORT -- FIELD KEY CHART

***** HEADING = 'MEGABYTES TRANSFERRED' ***** **1**

```

+-----+
| MEGABYTES |
| TRANSFERRED |
+-----+
      MMMMMMMM_ (SIZE OF EACH DATA SET IN MB)
FIELD----->I

```

FIELD	VALUES	MEANING
#1----	I	- DATA SET WAS PREVIOUSLY COMPACTED USING IDRC (IMPROVED DATA RECORDING CAPABILITY) ON THE DEVICE IT CURRENTLY RESIDES ON -- THIS WILL BE CONSIDERED WHEN CALCULATING THE EFFECT OF THE TARGET DEVICE ('3490E')

***** HEADING = 'CURRENT STAT-VOLS' ***** **2**

```

+-----+
| CURRENT |
| STAT-VOLS |
+-----+
      VVVVVV.VV_ (# OF CURRENT STATISTICAL VOLUMES FOR THE DATA SET)
FIELD----->I

```

FIELD	VALUES	MEANING
#1----	*	- THE VOLUMES ASSOCIATED WITH THIS DATA SET NO LONGER EXIST, AS THEY HAVE BEEN RETURNED TO SCRATCH FOR REUSE. THE VALUE SHOWN IS THE ORIGINAL NUMBER OF STATISTICAL VOLUMES FOR THE DATA SET.

Figure 52. The IDRC Report Legend

1 The MEGABYTES TRANSFERRED column shows the size of the most current version of the data set.

2 The CURRENT STAT-VOLS column shows the current statistical volumes for the data set. See “Statistical Volumes” on page 31 for more information about statistical volumes.

Example of the IDRC Report

Figure 53 on page 124 shows the possible effects of IDRC compaction, 3490E enhanced recording format, and the Enhanced Capacity Cartridge System Tape on large data sets.

GFTASRT3 -- IDRC REPORT -- IDRC FACTOR = 3.0 -- TAPEDEV = 3490E -- TAPELEN = 2 **3**

DATA SET NAME	DSN #	S/M FILE	TAPE DEVICE	MEGABYTES TRANSFERRED	# CURRENT STAT -VOLS	# CURRENT STAT -MNTS	3490E	3490E
							STAT -VOLS	STAT -MNTS
BAGI.FG571.BACKUP.G1028V00	1	SNGL	348X	2057I	4.00	28.00	1.00	7.00
BCAP.CAP.UNLD.DDACCAPR	2	MULT	348X	2160I	2.00*	4.00	0.49*	0.98
BCAP.CAP.UNL2.DDACCAPR	3	SNGL	348X	2143I	2.00*	5.00	1.00*	2.50
DKED.CIS.CICUSAPR.UNLOAD.G1061V00	4	SNGL	348X	1569I	3.00	6.00	1.00	2.00
DKED.CIS.CICUSAPR.UNLOAD.G1062V00	5	SNGL	348X	1573I	3.00	6.00	1.00	2.00
DKED.CIS.CICUSAPR.UNLOAD.G1063V00	6	SNGL	348X	1573I	3.00	6.00	1.00	2.00
IBMC.BMCS.BACKUP.MONTHS.G1686V00	7	MULT	348X	1840I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1687V00	8	MULT	348X	1844I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1688V00	9	MULT	348X	1848I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1689V00	10	MULT	348X	1851I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1690V00	11	MULT	348X	1860I	2.00*	2.00	0.48*	0.48
IBMC.BMCS.BACKUP.MONTHS.G1691V00	12	MULT	348X	1913I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1692V00	13	MULT	348X	1931I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1693V00	14	MULT	348X	1750I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1694V00	15	MULT	348X	1759I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1695V00	16	MULT	348X	1771I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1696V00	17	MULT	348X	1781I	2.00*	2.00	0.48*	0.48
IBMC.BMCS.BACKUP.MONTHS.G1697V00	18	MULT	348X	1787I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1698V00	19	MULT	348X	1793I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1699V00	20	MULT	348X	1798I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1700V00	21	MULT	348X	1807I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1701V00	22	MULT	348X	1812I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1702V00	23	MULT	348X	1812I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1703V00	24	MULT	348X	1816I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1704V00	25	MULT	348X	1816I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1705V00	26	MULT	348X	1821I	2.00*	2.00	0.48*	0.48
IBMC.BMCS.BACKUP.MONTHS.G1706V00	27	MULT	348X	1825I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1707V00	28	MULT	348X	1832I	2.00*	2.00	0.49*	0.49
IBMC.BMCS.BACKUP.MONTHS.G1708V00	29	MULT	348X	1837I	1.17	2.00	0.29	0.49
IBMC.BMCS.BACKUP.MONTHS.G1709V00	30	MULT	348X	1841I	1.50	2.00	0.37	0.49
IBMC.BMCS.BACKUP.MONTHS.G1710V00	31	MULT	348X	1844I	1.67	2.00	0.41	0.49
IBMC.BMCS.BACKUP.MONTHS.G1711V00	32	MULT	348X	1848I	1.00	2.00	0.25	0.50
IBMC.BMCS.HISTM.CICCSU.G1108V00	33	MULT	348X	1636I	1.07	1.07	0.27	0.27
IBMC.BMCS.HISTM.CICCSU.G1109V00	34	MULT	348X	1694I	1.07	1.07	0.26	0.26
IBMC.BMCS.MBACKUP.MONTHS.G1105V00	35	MULT	348X	1738I	1.70	2.00	0.42	0.49
DEPTAB.B.NEW.ACCOUNT	36	SNGL	348X	2185I	2.00*	4.00	1.00*	2.00
DEPTAB.D.OLD.ACCTSEGS	37	SNGL	348X	2192I	2.00*	7.00	1.00*	3.50
SYS2.HELD.OUTPUT.JUN16	38	SNGL	348X	5199I	2.00*	4.00	1.00*	2.00
SYS2.HELD.OUTPUT.JUN22	39	SNGL	348X	2152I	1.00*	2.00	1.00*	2.00
SYS2.HELD.OUTPUT.JUN24	40	SNGL	348X	3968I	2.00	4.00	1.00	2.00
SYS2.NONHELD.OUTPUT.JUN24	41	SNGL	348X	4960I	2.00	4.00	1.00	2.00
T O T A L S				84236	26.18	136.15	8.27	43.74

GFTAVMA PROGRAM COMPLETED AT 10:59:56 A.M. ON 06/03/1992 -- RETURN CODE = 0

Figure 53. The IDRC Report

3 The target tape device is controlled by the **TAPEDEV** keyword, which was specified as **3490E**, and the **TAPELEN** keyword, which was specified as 2. **IDRCFACTOR** was specified as 3.

4 In the **S/M FILE** column, **SNGL** or **MULT** indicates whether a data set is single-file or multifile. See “Single/Multi Volume/File Data Sets” on page 29 for more details.

5 The **MEGABYTES TRANSFERRED** column is the size, in megabytes, of the current version of a data set. This size is used to calculate the number of target (**TAPEDEV**) volumes required. The specified **IDRCFACTOR** is used to determine the effective capacity of the target device type. Because **GFTAVMA** cannot recognize data sets that have been software compacted, **GFTAVMA** does not account for software compaction when calculating the effective device capacity.

6 An asterisk (*) in the **# CURRENT STAT-VOLS** column shows that the volume has been reused. If the volumes associated with a data set no longer exist, there is an ‘*’ following the current statistical volumes value. In this case, the current value for statistical volumes is the *original* number of statistical volumes associated with the data set.

GFTAVMA does not count a volume more than once. To avoid double counting, the reused volumes are not accumulated in the totals for the IDRC Report. Therefore, the columns for current (**6**) and projected (**7**) statistical volumes do not add up to their totals.

7 The projected numbers of statistical volumes and statistical mounts result from the combination of the **IDRCFACTOR**, **TAPEDEV**, and **TAPELEN** keywords. An asterisk (*) in the 3490E STAT-VOLS column shows that the volume has been reused.

The Management Class Report

The Management Class Report is a simulation report. The simulation reports show how you can change the tape usage at your installation using tape mount management.

Use the keyword **REPORT(MGMTCLAS)** or **REPORT(MC)** to generate the Management Class Report. Using the **DATE** and **TIME** keywords for the Management Class Report impacts the accuracy of the report.

Note: Do not specify **LEVEL0AGE** or **LEVEL1AGE** keywords for this report because its purpose is to determine the optimum values for these keywords.

In this example, this keyword was specified:

```
REPORT(MGMTCLAS)
```

Since no other keywords were specified,

- **BLKSIZE** defaulted to 32 760
- **COMPACTION** defaulted to 50%
- **DASDDEV** defaulted to 3390-3
- **IDRCFACTOR** defaulted to 4
- **LARGE** defaulted to 600MB
- **TAPEDEV** defaulted to 3590
- **TAPELEN** defaulted to 1

If you want to model another DASD device, specify the **DASDDEV** keyword.

Keywords that Affect the Management Class Report

The results of the GFTAVMA reports rely on the keywords you choose. Therefore, choose your keywords carefully and be aware of the defaults and effects of the keywords on these reports.

All filtering input keywords and the **SPLIT** keyword affect these reports. See Chapter 5, "Using GFTAVMA Keywords to Filter Your Input," on page 63 and "SPLIT Keyword" on page 151 for the syntax and functions of these keywords.

A subset of the tailoring output keywords specified for a GFTAVMA run also affects the Management Class Report. See Chapter 7, "Using GFTAVMA Keywords to Tailor Your Output," on page 131 for more information about the syntax and functions of these tailoring output keywords:

- **BLKSIZE**
- **BPROGRAM**
- **CLASSIFY**
- **COMPACTION**
- **DASDDEV**
- **FORCECAP**

Management Class Report

- IDRCFACTOR
- LARGE
- RPROGRAM
- TAPEDEV
- TAPELEN

See Appendix A, “A Quick Reference to Volume Mount Analyzer Keywords,” on page 157 and Appendix B, “Volume Mount Analyzer Keyword Syntax,” on page 159 for more information about GFTAVMA keywords.

Description

The Management Class Report shows which management class values would be most effective in managing data in the storage hierarchy for your installation. The **LEVEL0AGE** and **LEVEL1AGE** values are the number of days since a data set's last reference. Therefore, when you specify these values, you tell GFTAVMA to model the primary volumes and migration level 1 volumes so that data sets migrate based on the management classes PRIMARY DAYS NON-USAGE and LEVEL 1 DAYS NON-USAGE. GFTAVMA does not model TEMP, LARGE, or BKUP management classes.

GFTAVMA simulates the data set migration and recall patterns that would occur during the sample period based on various combinations of the **LEVEL0AGE** and **LEVEL1AGE** keywords.

Use of the Management Class Report

The Management Class Report helps you determine the cost of DASD storage depending on the choice of **LEVEL0AGE** and **LEVEL1AGE** keyword values. For example, if you create a data set and reference it every five days, an **L0AGE** value of 1 and an **L1AGE** value of 6 will be optimum if you want to avoid any migration to migration level 2 tape. On the other hand, you might save more DASD space by setting **L0AGE** to 0 and **L1AGE** to 0. In doing so, the data sets will migrate directly from primary DASD to migration level 2 tape.

Since the Management Class Report lists the combinations for **LEVEL0AGE** and **LEVEL1AGE**, run the Management Class Report when you run the Estimate Report to determine the optimum values for **LEVEL0AGE** and **LEVEL1AGE**. Based on your 30-day SMF sample, GFTAVMA determines your data set usage patterns and reports on the optimum management class values for the active data sets.

Once you have selected a **L0AGE** and **L1AGE** combination for a given set of data, run the Estimate Report again, specifying **CLASSIFY(ACTIVE)** and the optimum **L0AGE** and **L1AGE** values.

See *z/OS DFSMS: Implementing System-Managed Storage* and *MVS/ESA SML: Managing Data* for details about using the Management Class and Estimate Reports together.

Assumptions for the Management Class Report

When generating the Management Class Report, GFTAVMA assumes you specified the migration level 1 compaction value that you feel you can achieve using the DFSMSHsm level 1 compaction option (default is 50% compaction, which means level 1 is reduced by 50% from primary storage).

Example of the Management Class Report

The Management Class Report shows the combinations of **L0AGE** and **L1AGE** that would use DASD most cost effectively. For each combination, it also shows the number of recall mounts that would still occur. Recall mounts are mounts that result from data sets that would be called back from migration level 2 tape to primary DASD.

Figure 54 shows *part* of the Management Class Report. This run produced a total of 528 combinations of the **LEVEL0AGE** and **LEVEL1AGE** keywords.

GFTASRT3 -- REPORT(MGMTCLAS)

ORDER	L0AGE	L1AGE	L2-MNTS	RECALLS	TOT TRKS	L0-TRKS	L1-TRKS	L2-MB	
1	1	1	3384	3384	103622	103622	0	1689	1
2	0	1	3384	3384	83	0	83	1	2
3	0	0	3384	3384	0	0	0	2993	
4	1	2	1200	3384	103656	103622	34	1	
5	0	2	1200	3384	114	0	114	1	
6	2	2	1200	1200	145856	145856	0	1196	
7	1	3	900	3384	103678	103622	56	0	
8	0	3	900	3384	132	0	132	0	
9	2	3	900	1200	145881	145856	25	0	
10	3	3	900	900	171418	171418	0	791	
11	1	4	485	3384	103687	103622	65	0	
12	0	4	485	3384	139	0	139	0	
13	2	4	485	1200	145892	145856	36	0	3 <====
14	3	4	485	900	171431	171418	13	0	
15	4	4	485	485	185339	185339	0	560	
16	1	5	351	3384	103697	103622	75	0	
17	0	5	351	3384	148	0	148	0	
18	2	5	351	1200	145902	145856	46	0	
19	3	5	351	900	171441	171418	23	0	
20	4	5	351	485	185349	185339	10	0	

Figure 54. Partial Management Class Report

Note: Each row of this report represents a unique management class.

Management Class Report

1 The first combination of **L0AGE** and **L1AGE** for this run is 1 and 1. These values leave the data sets on primary storage for at least 24 hours before the data sets migrate to level 2 tape. Because the value for **L1AGE** is the same as the value for **L0AGE**, data sets would bypass migration level 1 storage. This combination would result in 3384 mounts and recalls.

2 The second combination of **L0AGE** and **L1AGE** for this run is 0 and 1. These values allow the data sets on primary DASD to migrate to migration level 1 DASD when the next migration occurs, which is less than 24 hours. Then the data sets would remain on migration level 1 for at least 24 hours before migrating to level 2 tape. This combination would result in 3384 mounts and recalls.

3 The thirteenth combination of **L0AGE** and **L1AGE** for this run is 2 and 4. This combination significantly reduces migration level 2 mounts from 3384 to 485 and reduces recalls from 3384 to 1200, and the combination of 2 and 4 would be good values to specify for **L0AGE** and **L1AGE** for a study at this installation.

Understanding the Columns of the Report

The columns of the report are defined as follows:

ORDER:

The report is sorted in descending order:

1. Most mounts incurred
2. Most recalls incurred (from level 1 and level 2)
3. Most DASD tracks used

L0AGE:

The setting for primary days for this management class

L1AGE:

The setting for migration level 1 days for this management class

L2-MNTS:

The number of migration level 2 mounts that would be incurred for recalls

RECALLS:

Total number of recalls from migration level 1 or migration level 2 to primary DASD

TOT TRKS:

The total number of DASD tracks (*weighted average* over time) that would be needed

L0-TRKS:

The number of required tracks of primary DASD (*weighted average* over time). This value is not the total number of tracks containing data; it is the number of tracks needed for migration and recall activity.

L1-TRKS:

The number of required tracks of migration level 1 DASD (*weighted average* over time). This value is not the total of tracks containing data; it is the number of tracks needed for migration activity, not recall activity. Nothing is ever recalled to migration level 1, only to primary storage.

L2-MB:

Number of megabytes of data (*weighted average*) that would reside on migration level 2 tape because of migration activity.

Understanding the Values in the Management Class Report

Because GFTAVMA calculates DASD tracks as a weighted average over time, all of the DASD track values in Figure 54 on page 127 represent a weighted average based on the number of days the data set resided on the various levels of storage.

For example, a data set with 500 tracks and a size of 15MB resides on migration level 2. If it resided on migration level 2 for two-thirds of the sample, the Management Class Report would report its size as 10MB. The data set would also have spent one-third of its life on DASD using one-third of 500 tracks, which is 167 tracks of primary DASD.

Management Class Report

Chapter 7. Using GFTAVMA Keywords to Tailor Your Output

This chapter describes the GFTAVMA keywords you use to tailor the output reports of the volume mount analyzer. Unlike the keywords that filter your input, these keywords affect the output phase of processing by reducing the physical size of the reports, creating specific reports, and specifying the environment to be modelled in the detailed reports.

As pointed out earlier, all GFTAVMA keywords are optional. The VMACNTL DD statement points to a sequential data set with 80-byte records that contain the GFTAVMA keywords. Normally, this is a DD * data set. See the GFTAVMAP JCL member in “Submitting the GFTAVMA JCL” on page 48.

The keywords in this chapter are arranged in alphabetical order. Sub-keywords, if any, are described with their corresponding keyword.

Keyword Rules

When using keywords to tailor your output, you should:

- Start or end keywords in any position from 1 to 71. Positions 72 through 80 are ignored
- Separate keywords with either blanks or commas
- Separate sub-keywords with commas
- Separate items in a list with commas
- Put keywords in any order
- Use correct number of parentheses

If a keyword spans more than one record, you must indicate continuation by using continuation characters in the last text character of the keyword. Blanks can follow a continuation character in a record. Valid continuation characters for GFTAVMA are a comma and the left and right parenthesis.

Comments are surrounded by the delimiters /* and */. A comment start delimiter, /*, cannot be in column one. Examples:

```
TIME(0,12) /* This is a comment */
TIME(0,    /* The last character, the comma, indicates continuation */
  12)    /* part two of the continued keyword */
```

Summary of Keywords That Tailor Your Output Reports

Table 6 summarizes the output report keywords.

See Appendix A, “A Quick Reference to Volume Mount Analyzer Keywords,” on page 157 for keyword abbreviations and defaults. See Appendix B, “Volume Mount Analyzer Keyword Syntax,” on page 159 for keyword syntax.

Table 6. Keywords that Tailor Your Output Reports

Keyword	Reports Affected	Page Number
BLKSIZE	Data Set, GBMAX, Top, Volume, Estimate, IDRC, Management Class	132
BPROGRAM	Estimate, Management Class	133
CHART	None	134

Table 6. Keywords that Tailor Your Output Reports (continued)

Keyword	Reports Affected	Page Number
CLASSIFY	Estimate, Management Class	135
COMPACTION	Estimate, Management Class	136
CUTOFF	Data Set, Volume	138
DASDDEV	Data Set, Volume, Estimate, Management Class	138
DSORT	Data Set	139
FORCECAP	Data Set, GBMAX, Volume, Estimate, IDRC, Management Class	140
NOFORCECAP	Data Set, GBMAX, Volume, Estimate, IDRC, Management Class	140
FSPACE	Estimate	141
GDG ¹	Data Set, GFTASRT3 Summary Report, Estimate, IDRC	142
IDRCFACTOR	Estimate, IDRC, Management Class	143
LARGE	Estimate, IDRC, Management Class	144
LEVEL0AGE ²	Estimate	145
LEVEL1AGE ²	Estimate	146
LINES	All	147
LIFSPACE	Estimate	147
REPORT	All detailed reports	148
RPROGRAM	Data Set, GBMAX, Volume, Estimate, IDRC, Management Class	150
SPLIT	All	151
TAPEDEV	Estimate, IDRC, Management Class	152
TAPELEN	Estimate, IDRC, Management Class	153
VSORT	Volume	154

Notes:

1. Do not specify **GDG(GROUP)** when you request the Estimate Report or the IDRC Report. GFTAVMA does not allow you to request **GDG(GROUP)** with **REPORT(EST)** or **REPORT(IDRC)**. Specify it *only* when you request the Data Set Report.
2. Do not specify **LEVEL0AGE** or **LEVEL1AGE** when you request the Management Class Report because the Management Class Report shows the most effective values for these keywords.

BLKSIZE Keyword

Use **BLKSIZE** to supply a block size value for GFTAVMA to use when the SMF records indicate a block size of zero.

GFTAVMA uses the block size supplied in the SMF records, unless that value is zero. In that case, GFTAVMA uses the value you specified with the **BLKSIZE** keyword.

Abbreviation: BLKS

Syntax

Figure 55 shows the syntax for the **BLKSIZE** keyword.

[BLKSIZE (nnnnn)]

Figure 55. *BLKSIZE* Keyword Syntax

nnnnn

Represents the number of bytes you want to specify for block size.

SMF data does not supply the total amount of data transferred. Therefore, GFTAVMA must calculate this value from the block size and the number of blocks in the written fields of the SMF type 14/15 records.

Default

If you do not specify the **BLKSIZE** keyword, GFTAVMA uses its default of 32 760.

BLKSIZE Example

If you specify:

```
BLKSIZE (27998)
```

you have chosen a good starting point when estimating half-track blocking. This is because you can put two 27998 blocks on a 3390 track. You can specify a larger block size, but it would not be as efficient.

Reports Affected

BLKSIZE affects:

- The Data Set Report (see “The Data Set Report” on page 82 for more information)
- The GBMAX Report (see “The Maximum Gigabyte Report” on page 88 for more information)
- The Top Report (see “The Top Report” on page 92 for more information)
- The Volume Report (see “The Volume Report” on page 101 for more information)
- The Estimate Report (see “The Estimate Report” on page 106 for more information)
- The IDRC Report (see “The Improved Data Recording Capability Report” on page 120 for more information)
- The Management Class Report (see “The Management Class Report” on page 125 for more information)

BPROGRAM Keyword

When you use **BPROGRAM** to designate a list of program names as backup programs, GFTAVMA considers any data set that one of these programs creates to be a backup data set if the data set’s usage pattern is backup.

If you do not agree with the data set category that the volume mount analyzer assigns to data sets, you can change the category of selected data sets using the **BPROGRAM** keyword. See “Categories of Tape Mount Management Data Sets” on page 38 for more information on the logic GFTAVMA uses to categorize data sets as either active or backup.

BPROGRAM

Abbreviation: BPGM

Syntax

Figure 56 shows the syntax for the **BPROGRAM** keyword.

```
[BPROGRAM (INCLUDE (backup program include list),  
          EXCLUDE (backup program exclude list))]
```

Figure 56. **BPROGRAM** Keyword Syntax

INCLUDE (*backup program include list*)

Specifies a list of backup programs to be included in the analysis.

EXCLUDE (*backup program exclude list*)

Specifies a list of backup programs to be excluded from the analysis.

Default

If you do not specify the **BPROGRAM** keyword, GFTAVMA uses the data set usage types assigned during GFTAXTR processing. These types are based on usage patterns determined during the time period of the study. See “Data Set Usage Types” on page 36 for more details.

Relationship to Other Keywords

Do not use the **CLASSIFY** and **BPROGRAM** keywords together in the same run. This is because while **CLASSIFY** sets the entire sample to one tape mount management category, **BPROGRAM** categorizes selected data sets.

If you do not specify the **CLASSIFY** keyword, GFTAVMA considers a program designated by the **BPROGRAM** keyword to be a backup program. For GFTAVMA to categorize a data set as the tape mount management category of backup, the data set must also have a usage pattern of backup (zero reads and one or more writes).

Reports Affected

BPROGRAM affects the Estimate Report and the Management Class Report.

See “The Estimate Report” on page 106 and “The Management Class Report” on page 125 for more details.

CHART Keyword

Use **CHART** to generate a sequential data set for the run. The sequential data set contains formatted data that you can use as input to your graphics package.

Abbreviation: None

Syntax

Figure 57 on page 135 shows the syntax for the **CHART** keyword.

[CHART]

Figure 57. CHART Keyword Syntax

CHART has no sub-keywords.

Additional JCL Needed for CHART

The **CHART** keyword creates a sequential data set in the data set referenced by the VMACHART DD statement.

In the GFTAVMAP JCL member, the VMACHART DD is allocated as a SYSOUT file. To override this specification, insert the following DD statement in your GFTAVMA execution input stream following the EXEC GFTAVMA statement:

```
//VMA.VMACHART DD DSN=your.data.set, DISP=(NEW,CATLG),
//   DCB=(LRECL=80,RECFM=FB), SPACE=(TRK, (1,1))
```

This statement directs the **CHART** output to a sequential data set that you can use to create charts using the graphics package of your choice. See “Submitting the GFTAVMA JCL” on page 48 for the complete JCL member.

Default

If you do not specify the **CHART** keyword, GFTAVMA does not generate a sequential data set for the run.

CHART Example

See Appendix C, “Example of Chart Output,” on page 161 for examples of the output **CHART** generates.

Reports Affected

CHART does not affect any of the GFTAVMA reports.

CLASSIFY Keyword

Use **CLASSIFY** to designate *all* data sets included in the GFTAVMA run as either active or backup.

You can classify all data sets as either active or backup to determine the effects of the values for **L0AGE** and **L1AGE**. Classify all data sets as backup to simulate migrating all data sets at the earliest possible time (at least once a day). Classify all data sets as active to simulate migrating all data sets based on the **L0AGE** and **L1AGE** values.

The **CLASSIFY** keyword overrides the tape mount management categories GFTAVMA assigned. See “Categories of Tape Mount Management Data Sets” on page 38 for more information on the logic GFTAVMA uses to categorize data sets.

Abbreviation: CLASS

Syntax

Figure 58 on page 136 shows the syntax for the **CLASSIFY** keyword.

CLASSIFY

```
[CLASSIFY (ACTIVE | BACKUP)]
```

Figure 58. CLASSIFY Keyword Syntax

ACTIVE

Designates all data sets in the run (other than HSM, LARGE, and TEMP) as active.

BACKUP

Designates all data sets in the run (other than HSM, LARGE, and TEMP) as backup.

You can specify only one sub-keyword on the **CLASSIFY** keyword, either active or backup, but not both.

Default

If you do not specify the **CLASSIFY** keyword, GFTAVMA uses the tape mount management categories described in “Categories of Tape Mount Management Data Sets” on page 38.

Relationship to Other Keywords

Do not specify the **CLASSIFY** keyword in the same run as the **BPROGRAM** keyword. This is because **CLASSIFY** sets the entire sample to one usage class, while **BPROGRAM** classifies selected data sets.

Specify **CLASSIFY** only when you request the Estimate Report or the Management Class Report. When *modelling* the effects of DFSMS, GFTAVMA uses the data set categories to model active data in the DFSMSHsm storage hierarchy. If the data set is categorized as backup, GFTAVMA models the data set as backup in the DFSMSHsm hierarchy.

Restrictions

The **CLASSIFY** keyword does not change the designations for:

- DFSMSHsm single file format data sets
- Temporary data sets
- Large data sets

Reports Affected

CLASSIFY affects the Estimate Report and the Management Class Report.

See “The Estimate Report” on page 106 and “The Management Class Report” on page 125 for more details.

COMPACTION Keyword

Use **COMPACTION** to specify the compaction percentage. GFTAVMA uses the same compaction percentage to model the effects of DFSMSHsm software compaction when calculating the required number of migration level 1 DASD volumes for the Estimate Report and the Management Class Report.

Abbreviation: COMP

Syntax

Figure 59 shows the syntax for the **COMPACTION** keyword.

[COMPACTION (<i>nn</i>)]

Figure 59. **COMPACTION** Keyword Syntax

nn Represents the percentage of compaction achieved. This value indicates that all data sets in the sample can achieve *nn*% compaction when DFSMSHsm moves the data sets from primary to migration level 1 DASD.

Default

If you do not specify the **COMPACTION** keyword, GFTAVMA uses the default of 50%.

Restrictions

COMPACTION only affects active data sets. GFTAVMA models backup data sets as migrating directly to migration level 2 and to whatever device you specify with the **TAPEDEV** and **TAPELEN** keywords.

DFSMSHsm compaction is taken into account for data sets that GFTAVMA determines to belong on migration level 2 tape based on the length of time since they have been referenced. These data sets would have first resided on migration level 1 for some period of time. For these data sets, GFTAVMA applies the *nn* percentage compaction that would have been achieved during migration to migration level 1 and then assumes an additional 15% IDRC compaction during the level 1 to level 2 migration.

Data sets that migrate directly from primary storage to migration level 2 do not require DFSMSHsm software compaction if IDRC is available to the tape unit controllers (3480/3490). This saves processor and IDRC compaction at the bit level and is more efficient than the software compaction at the byte level. Data sets that are already software compacted can usually only be IDRC compacted an additional 10%.

Therefore, when a 600MB data set with a specified **IDRCFACTOR** of 3.0 migrates from primary DASD to migration level 2 tape, the data set is compacted to 200MB. If the same 600MB data set has **COMPACTION(50)** and **IDRCFACTOR(3.0)**, it is compacted 50% to 300MB when it migrates to migration level 1. When it migrates to migration level 2, it is reduced an additional 45MB to 255MB. Thus, the IDRC compacted data set achieves a total of 200MB, while the software compacted and IDRC compacted data set achieves only a total of 255MB.

Therefore, you must consider the effects of using DFSMSHsm's software compaction with IDRC compaction. A data set that has been software compacted cannot achieve full IDRC compaction.

COMPACTION Example

If you specify:

```
COMPACTION(90)
```

data sets on migration level 1 volumes will be reduced to 90% of their original size.

Reports Affected

COMPACTION applies only to DFSMSHsm migration level 1 volumes and affects the Estimate Report and the Management Class Report.

See “The Estimate Report” on page 106 and “The Management Class Report” on page 125 for more details.

CUTOFF Keyword

Use CUTOFF to stop the printing of a Data Set or Volume Report after a specified number of gigabytes of data has been reached for that report. This keyword has no effect on the values shown in any of the volume mount analyzer reports.

Use this keyword to save paper or cut down on the size of the output listing.

Abbreviation: CUT

Syntax

Figure 60 shows the syntax for the CUTOFF keyword.

[CUTOFF (nnnnnnnnnn)]

Figure 60. CUTOFF Keyword Syntax

nnnnnnnnnn

Specifies the number of gigabytes printed in a report.

The value specified can be from 1 to 2147483647.

Default

If you do not specify the CUTOFF keyword, GFTAVMA uses a cutoff of infinity.

CUTOFF Example

If you specify:

CUTOFF (85)

your report will stop printing when the cumulative gigabytes exceed 85GB.

Reports Affected

CUTOFF affects the Data Set Report and the Volume Report.

See “The Data Set Report” on page 82 and “The Volume Report” on page 101 for more details.

DASDDEV Keyword

Use DASDDEV to control the target DASD device that GFTAVMA uses for all calculations involving DASD space and DASD track sizes in the GFTAVMA reports. All DASD/TRK values reflect the selected device type.

Abbreviation: DASD

Syntax

Figure 61 shows the syntax for the **DASDDEV** keyword.

```
[DASDDEV (3380S | 3380D | 3380J | 3380E | 3380K
          | 3390-1 | 3390-2 | 3390-3 | 3390-9 | 9345-1 | 9345-2)]
```

Figure 61. DASDDEV Keyword Syntax

Select one of these DASD devices as the target device to which GFTAVMA should model data.

Default

The default for **DASDDEV** is 3390-3.

DASDDEV Example

For each data set, the Data Set Report shows the number of DASD tracks that would be required to hold the data. The value used for **DASDDEV** controls which DASD track size is used for this calculation.

The actual track calculation is based on both the **DASDDEV** specification and data set attributes (such as LRECL, BLKSIZE, and RECFM).

Reports Affected

DASDDEV affects:

- The Data Set Report (see “The Data Set Report” on page 82 for more information)
- The Volume Report (see “The Volume Report” on page 101 for more information)
- The Estimate Report (see “The Estimate Report” on page 106 for more information)
- The Management Class Report (See “The Management Class Report” on page 125 for more information)

DSORT Keyword

Use **DSORT** to specify how the list of data set names in the Data Set Report is sorted. When GFTAVMA sorts on sub-keywords other than data set names, there is a secondary sort performed by data set name within the major sort category.

Abbreviation: DSR

Syntax

Figure 62 shows the syntax for the **DSORT** keyword.

```
[DSORT (DATASET
        | JOBNAME | PROGRAM | SIZE | EXPDT)]
```

Figure 62. DSORT Keyword Syntax

DSORT

DATASET (DSN)

Prints the list of data sets in ascending order by data set name. This is the default and the most useful option for helping you analyze random tape mounts.

JOBNAME (JOB)

Prints the list of data sets in ascending order by job name. This option can be useful for analyzing all data sets associated with a particular job name, such as all data sets related to DFSMSHsm processing.

PROGRAM (PGM)

Prints the list of data sets in ascending order by program name. This option can be useful for analyzing certain data sets associated with a particular program.

SIZE (SIZ)

Prints the list of data sets in ascending order by data set size.

EXPDT (EXP)

Prints the list of data sets in ascending order by expiration date.

Default

The default for **DSORT** is DATASET.

Reports Affected

DSORT only affects the Data Set Report. See “The Data Set Report” on page 82 for more details.

FORCECAP/NOFORCECAP Keywords

Use **FORCECAP** to set an upper limit on the amount of data that GFTAVMA has determined to be on a single volume based on the relative device capacity of that volume.

Note: You should always run with **FORCECAP** because it shows the block size of the data as close as possible to its actual total.

Abbreviations: **FCAP** and **NOFCAP**

Syntax

Figure 63 shows the syntax for the **FORCECAP** keyword.

[<u>FORCECAP</u> NOFORCECAP]

Figure 63. **FORCECAP/NOFORCECAP** Keyword Syntax

FORCECAP

Sets a limit on the amount of data calculated to be on a volume based on the capacity of the device.

NOFORCECAP

Does not set a limit on the amount of data calculated to be on a volume. As a result, GFTAVMA might calculate a larger size than the capacity of the device.

Note: Run with **NOFORCECAP** only if you wish to see how high the calculations can be.

The Importance of FORCECAP

Using block size and number of blocks written, GFTAVMA calculates the total amount of data transferred. For variable blocked and undefined records, SMF supplies the user-stated block size found in the JFCB. If the block size is zero, GFTAVMA uses the **BLKSIZE** keyword value, which defaults to 32 760. SMF does not indicate the total amount of data written or the average block size used.

Therefore, in the case of records whose format is variable length or undefined, GFTAVMA's calculation of the total amount of data transferred is usually somewhat higher than the actual amount of data that was written. This over-calculation is evident when GFTAVMA calculates that a single volume contains well in excess of its maximum capacity.

GFTAVMA attempts to compensate for this over-calculation by setting a maximum capacity (**FORCECAP**) of data per volume: 416MB for 3490E cartridges, 208MB for 3490 and 3480 cartridges, and 169MB for 3420 tape volumes. For data sets that have been IDRC compacted, GFTAVMA adjusts volumes written with IDRC to the specified **IDRCFACTOR** keyword value multiplied by the device capacity ($IDRCFACTOR \times devcap$).

Default

FORCECAP is the default.

If you specify **NOFORCECAP**, GFTAVMA uses the value calculated by block size multiplied by the number of blocks, which can exceed the maximum capacity of the volume.

Reports Affected

FORCECAP/NOFORCECAP affects:

- The Data Set Report (see "The Data Set Report" on page 82 for more information)
- The GBMAX Report (see "The Maximum Gigabyte Report" on page 88 for more information)
- The Volume Report (see "The Volume Report" on page 101 for more information)
- The Estimate Report (see "The Estimate Report" on page 106 for more information)
- The IDRC Report (see "The Improved Data Recording Capability Report" on page 120 for more information)
- The Management Class Report (see "The Management Class Report" on page 125 for more information)

FSPACE Keyword

Use **FSPACE** to specify the free space low threshold that GFTAVMA should model for DFSMSHsm primary storage volumes when calculating the required number of volumes for the tape mount management DASD buffer.

Abbreviation: FSP

Syntax

Figure 64 on page 142 shows the syntax for the **FSPACE** keyword.

FSPACE

```
[FSPACE (nn)]
```

Figure 64. FSPACE Keyword Syntax

nn Represents the percentage for the DFSMSShsm free space low threshold.

The actual free space you need for a tape mount management implementation should be based on a **REPORT(GBMAX)** run against *all* of the filters planned for ACS routines.

See *z/OS DFSMS: Implementing System-Managed Storage* for details about determining free space.

Default

The default for **FSPACE** is 5%.

Reports Affected

FSPACE only affects the Estimate Report. See “The Estimate Report” on page 106 for more details.

GDG Keyword

Use **GDG** to specify how GDSs are listed in the Data Set Report. Each GDS can be listed individually or listed collectively by its GDG.

Abbreviation: None

Syntax

Figure 65 shows the syntax for the **GDG** keyword.

```
[GDG (SHOW | GROUP)]
```

Figure 65. GDG Keyword Syntax

SHOW

Prints out each GDS separately on the Data Set Report. This option can be useful when you need an in-depth analysis of each GDS.

GROUP

Groups all the GDSs from a GDG base together and prints one line for all the data sets within a GDG. This option simplifies the Data Set Report to save paper.

Default

The default for **GDG** is SHOW, which displays the GDSs separately.

Reports Affected

Use the **GDG** keyword *only* for the Data Set Report. See “The Data Set Report” on page 82 for more details. When you specify **GDG(GROUP)**, run only the Data Set Report. **GDG** also affects the values in the GFTASRT3 Summary Report. See “GFTASRT3—Tape Data Set Statistics Report” on page 57 for more details.

Note: Do *not* specify **GDG(GROUP)** with the Estimate Report or the IDRC Report because GFTAVMA will end processing with message GFTA046E. See the messages in Appendix D, “Volume Mount Analyzer Messages,” on page 167.

IDRCFACTOR Keyword

Use **IDRCFACTOR** to specify the value that you believe is a reasonable projected IDRC compaction ratio for your environment. Since GFTAVMA does not examine your data directly, it has no way of projecting actual compaction ratios that you might have at your installation.

The IDRC feature available on the 3490 and some 3480s can significantly increase the effective cartridge data capacity. Using IDRC reduces the number of cartridges required for multivolume data sets and therefore the number of required tape mounts. The actual reduction in cartridges that can be achieved with IDRC compaction depends upon the data set’s block size and data characteristics.

GFTAVMA calculates the projected number of volumes and mounts in a simulated IDRC environment based on the size of the data set, the effective capacity of the target media with the **IDRCFACTOR** applied, and the **TAPEDEV** and **TAPELEN** specifications. **TAPEDEV** and **TAPELEN** control the selection of which target device type GFTAVMA would use.

Abbreviation: IDRC

Syntax

Figure 66 shows the syntax for the **IDRCFACTOR** keyword.

[IDRCFACTOR (nn.n)]

Figure 66. IDRCFACTOR Keyword Syntax

nn.n

Specifies a value between 1.0 and 99.9.

GFTAVMA and Compacted Data Sets

When GFTAVMA finds the flag in the SMF record that indicates a data set has been IDRC compacted, it uses the actual IDRC value for its calculations instead of the value you specified with the **IDRCFACTOR** keyword.

However, SMF records do not have a flag to indicate software compaction. Since GFTAVMA reports on the size of data transferred and cannot tell if a data set is already software compacted or compressed, the data set size could be larger than the data set transfer size. For data sets that have been software compacted, the volume mount analyzer does not know the actual size of the data sets or whether or not they have been compacted.

Default

The default for **IDRCFACTOR** is 4.0.

Relationship to Other Keywords

The **IDRCFACTOR** keyword provides a compaction ratio that is used to project savings from IDRC compaction when you specify the **REPORT(IDRC)** keyword.

IDRCFACTOR

See “The Improved Data Recording Capability Report” on page 120 for more information on the IDRC Report and the interaction between **IDRCFACTOR** and the **TAPEDEV** and **TAPELEN** keywords.

IDRCFACTOR estimates the native device capacity when **GFTAVMA** estimates the data set size and **FORCECAP** has been specified.

Reports Affected

IDRCFACTOR affects:

- The Estimate Report (see “The Estimate Report” on page 106 for more information)
- The IDRC Report (see “The Improved Data Recording Capability Report” on page 120 for more information)
- The Management Class Report (see “The Management Class Report” on page 125 for more information)

LARGE Keyword

Use **LARGE** to indicate to **GFTAVMA** the size of data sets that should be modelled as bypassing the tape mount management DASD buffer and going directly to tape.

Abbreviation: LRG

Syntax

Figure 67 shows the syntax for the **LARGE** keyword.

[LARGE (nnnnnnnnnn)]

Figure 67. **LARGE** Keyword Syntax

nnnnnnnnnn

Specifies a number between 0 and 2147483647MB.

If you are including an application that might have a few large data sets that you intend on intercepting regardless of size, which is a good technique, you should set **LARGE** to a value high enough to ignore these data sets in the **LARGE** category.

If you specify **LARGE (0)**, **GFTAVMA** evaluates all data sets as candidates for IDRC compaction and includes them in the IDRC Report.

Default

The default for **LARGE** is 600MB.

Relationship to Other Keywords

Data sets equal to or larger than the value specified for **LARGE** are modelled as being written directly to tape. The **TAPEDEV**, **TAPELEN**, and **IDRCFACTOR** keywords control the device type and IDRC compaction for the target device to which **GFTAVMA** models these data sets.

Data sets smaller than the value specified for **LARGE** are candidates for *system-managed storage* and migrate to migration level 1 or migration level 2 according to their active or backup classification and **LEVEL0AGE** and **LEVEL1AGE** values.

Reports Affected

LARGE affects:

- The Estimate (see “The Estimate Report” on page 106 for more information)
- The IDRC (see “The Improved Data Recording Capability Report” on page 120 for more information)
- The Management Class Reports (see “The Management Class Report” on page 125 for more information)

LEVEL0AGE Keyword

Use **LEVEL0AGE** to indicate to GFTAVMA the number of days that must elapse between *the date of the last reference* to a data set and the date the data set can be migrated from primary DASD to either migration level 1 DASD or migration level 2 tape.

LEVEL0AGE applies primarily to the *active* category of tape mount management data sets. GFTAVMA models backup data sets as going directly to tape at the end of the day unless you specify **L0AGE=0** and **L1AGE=0**. GFTAVMA models backup data sets as migrating to migration level 2 immediately after their use. Since temporary data sets expire on DASD, their **L0AGE** is always 1 and **L1AGE** does not apply because temporary data sets are deleted before they are eligible for migration to migration level 1.

Abbreviation: L0AGE

Syntax

Figure 68 shows the syntax for the **LEVEL0AGE** keyword.

[LEVEL0AGE (nnnn)]

Figure 68. LEVEL0AGE Keyword Syntax

nnnn

Represents a whole number of days between 0 and 9999.

Note: If you specify **L0AGE** as 0 and if it causes an immediate recall on the same day, a half day of residency is charged to both primary DASD and to the target level storage, either migration level 1 or migration level 2.

Default

The default for **LEVEL0AGE** is 1 day.

Relationship to Other Keywords

The value for **LEVEL1AGE** cannot be smaller than the value for **LEVEL0AGE**. If you specify a smaller value for **LEVEL1AGE** than for **LEVEL0AGE**, GFTAVMA changes the value of **LEVEL1AGE** so that it is equal to the value of **LEVEL0AGE**. If you specify **L0AGE=1** and **L1AGE=0**, GFTAVMA would change **L1AGE** to 1.

The **LEVEL0AGE** keyword models the effect of the management class attribute PRIMARY DAYS NON-USAGE.

LEVEL0AGE

Reports Affected

LEVEL0AGE only affects the Estimate Report. See “The Estimate Report” on page 106 for more details.

Do not specify LEVEL0AGE when you request the Management Class Report because this report determines the optimum values for LEVEL0AGE.

LEVEL1AGE Keyword

Use LEVEL1AGE to indicate to GFTAVMA the number of days that must elapse between *the date of the last reference* to a data set and the date the data set can be migrated from migration level 1 DASD to migration level 2 tape.

LEVEL1AGE applies primarily to the *active* category of tape mount management data sets. GFTAVMA models backup data sets as going directly to tape at the end of the day unless you specify L0AGE=0 and L1AGE=0. GFTAVMA models backup data sets as migrating to migration level 2 immediately after their use. Since temporary data sets expire on DASD, their L0AGE is always 1 and L1AGE does not apply because temporary data sets are deleted before they are eligible for migration to migration level 1.

Abbreviation: L1AGE

Syntax

Figure 69 shows the syntax for the LEVEL1AGE keyword.

[LEVEL1AGE (<i>nnnn</i>)]

Figure 69. LEVEL1AGE Keyword Syntax

nnnn

Represents a whole number of days between 0 and 9999.

Default

The default for LEVEL1AGE is the value you specify for LEVEL0AGE.

LEVEL1AGE defaults to the LEVEL0AGE value so that will GFTAVMA simulate bypassing migration level 1. Data sets are modelled as migrating directly from primary storage to migration level 2 tape.

Relationship to Other Keywords

The value you specify for LEVEL1AGE includes the LEVEL0AGE value. For example, assume LEVEL0AGE is set to 2, and LEVEL1AGE is set to 10. After two days of not being referenced on primary DASD, the data set would be eligible to migrate to migration level 1. After an additional 8 days on migration level 1, the data set would migrate to migration level 2 tape. Therefore, LEVEL1AGE is *not* equal to the number of days the data set resides on migration level 1.

The value for LEVEL1AGE cannot be smaller than the value for LEVEL0AGE. If you specify a smaller value for LEVEL1AGE than for LEVEL0AGE, GFTAVMA changes the value of LEVEL1AGE so that it is equal to the value of LEVEL0AGE. If you specify L0AGE=1 and L1AGE=0, GFTAVMA changes L1AGE to 1.

The **LEVEL1AGE** keyword models the effect of the management class attribute **LEVEL 1 DAYS NON-USAGE**.

Reports Affected

LEVEL1AGE only affects the Estimate Report. See “The Estimate Report” on page 106 for more details.

Do not specify **LEVEL1AGE** for the Management Class Report because this report determines the optimum values for **LEVEL1AGE**.

LINES Keyword

Use **LINES** to set the number of lines per page for the printed output report.

Abbreviation: None

Syntax

Figure 70 shows the syntax for the **LINES** keyword.

[LINES (nn)]

Figure 70. *LINES* Keyword Syntax

nn Represents the number of lines of output to print per page. The value must be a whole number between 9 and 99.

Default

The default for **LINES** is 60 lines per page.

Reports Affected

LINES affects all GFTAVMA reports.

L1FSPACE Keyword

Use **L1FSPACE** to specify the free space threshold that GFTAVMA should model for DFSMSShsm migration level 1 volumes when calculating the required number of migration level 1 volumes required for the tape mount management DASD buffer.

Abbreviation: L1FSP

Syntax

Figure 71 shows the syntax for the **L1FSPACE** keyword.

[L1FSPACE (nn)]

Figure 71. *L1FSPACE* Keyword Syntax

nn Represents the percentage of the volume that should be left as free space.

See *z/OS DFSMS: Implementing System-Managed Storage* for details about free space and DASD buffers.

Default

The default for **L1FSPACE** is 5%.

Reports Affected

L1FSPACE only affects the Estimate Report. See “The Estimate Report” on page 106 for more details.

REPORT Keyword

Use **REPORT** to specify the optional detailed reports that you want GFTAVMA to generate. You can specify any or all of the optional reports for a single GFTAVMA run.

Abbreviation: REP

Syntax

Figure 72 shows the syntax for the **REPORT** keyword.

```
[REPORT (DATASET,ESTIMATE,GBMAX,IDRC,  
MGMTCLAS,TOP,USAGE,VOLUME)]
```

Figure 72. **REPORT** Keyword Syntax

DATASET (DSN)

Produces a Data Set Report showing details about data set use. See “The Data Set Report” on page 82 for more information.

ESTIMATE (EST)

Produces an Estimate Report showing projected costs and savings that result from using tape mount management to manage a subset of data. See “The Estimate Report” on page 106 for more information.

GBMAX (GB)

Produces a Maximum Gigabyte Report showing the maximum number of DISP=NEW or DISP=MOD GB allocated every day by hour along with the maximum per day and hour. See “The Maximum Gigabyte Report” on page 88 for more information.

IDRC

Produces an IDRC Report showing the projected savings from using IDRC to compact large data sets. See “The Improved Data Recording Capability Report” on page 120 for more information.

MGMTCLAS (MC)

Produces a Management Class Report showing the recommended management class values to optimally manage a set of data. See “The Management Class Report” on page 125 for more information.

TOP

Produces a Top Report showing the most frequent high level qualifiers, low level qualifiers, programs, jobs, expiration dates, account codes or data set ages that are responsible for the top percentage of all mounts. The Top Report is the only report that has parameters. See “Parameters for the Top Report” on page 149. See “The Top Report” on page 92 for more information about the Top Report.

USAGE (USE)

Produces a Usage Report showing the maximum concurrent number of used tape drives and mounts incurred for each hour, as well as the totals and peak values for each hour and each day. See “The Usage Report” on page 95 for an example of this report.

VOLUME (VOL)

Produces a Volume Report showing all tape volumes. See “The Volume Report” on page 101 for more information.

Parameters for the Top Report

The TOP sub-keyword is the only sub-keyword that has parameters. The syntax of the TOP sub-keyword is:

```
REPORT(TOP(parm1,parm2...))
```

The valid parameters for the TOP sub-keyword are as follows:

HLQ	High level qualifier
LLQ	Low level qualifier
PROGRAM	Program name (abbreviation is PGM)
JOBNAME	Job name (abbreviation is JOB)
REFERENCE	Last reference age (abbreviation is REF)
EXPDT	Expiration date (abbreviation is EXP)
ACCOUNT	Accounting information, which is the first 8 bytes of the accounting field as specified on the application program’s JOB statement (abbreviation is ACCT). This is affected by the ACCTFLD keyword, which is described in “ACCTFLD Keyword” on page 15.
AGE	Number of days since last reference
PERCENT(nn)	The percentage of mounts that should be used as a cutoff point (abbreviation is PCT)
ALL	All the above parameters

The defaults for the TOP sub-keyword are REPORT (TOP(ALL,PERCENT(100))).

Default

REPORT does not have a default. If you do not specify REPORT, GFTAVMA only produces the summary reports.

Relationship to Other Keywords

Use the DSORT and VSORT keywords to sort the Data Set and Volume Reports.

REPORT Examples

If you specify:
REPORT(EST,TOP)

GFTAVMA will generate an Estimate Report and a Top Report. Since the Top Report has no specified parameters, GFTAVMA uses the defaults, TOP(ALL,PERCENT(100)).

REPORT

The Estimate Report and the Top Report are a recommended first run to help you get started.

If you specify:

```
REPORT (TOP (HLQ, PROGRAM, EXPDT, PERCENT (40)))
```

GFTAVMA will produce Top Reports showing the high level qualifiers, program names, and expiration dates that account for the top 40% of all mounts.

Reports Affected

REPORT creates all the GFTAVMA detailed reports. If you do not specify **REPORT** and its sub-keywords, GFTAVMA does not generate any detailed reports.

RPROGRAM Keyword

Use **RPROGRAM** to designate a list of program names whose associated data sets are reblockable.

Abbreviation: RPGM

Syntax

Figure 73 shows the syntax for the **RPROGRAM** keyword.

```
[RPROGRAM (INCLUDE (rprogram include list),  
EXCLUDE (rprogram exclude list))]
```

Figure 73. *RPROGRAM* Keyword Syntax

INCLUDE (*reblockable program include list*)

Specifies a list of program names whose data sets will be reblockable.

EXCLUDE (*reblockable program exclude list*)

Specifies a list of program names whose data sets will not be reblockable.

Many data sets currently on tape have block sizes that are efficient for tape, but inefficient for DASD. For example, 32 760 is a commonly used block size for tape data sets. However, 32 760 is greater than half a track of 3390 DASD and would result in poor track use. Fortunately, most programs have no restrictions on block size. In the DFSMS environment you can eliminate the use of hard coded block size and allow DFSMS to choose the optimum block size with system determined block size.

If the user's DD statement has a BLKSIZE value that exceeds 32 760 when writing, that data set cannot be redirected to DASD. If the system determined the block size and it exceeds 32 760, the system will determine a smaller block size if the data set is redirected to DASD.

GFTAVMA calculates tracks required using the device type specified by the **DASDDEV** keyword and the data set's block size and block count.

If you specify **RPROGRAM**, GFTAVMA models designated programs as reblockable. If GFTAVMA considers a program to be reblockable, any data sets created by that program and allocated to DASD are modelled as having optimum block sizes that use DASD tracks more efficiently.

Default

If you do not specify the **RPROGRAM** keyword, GFTAVMA does not consider any programs reblockable.

Reports Affected

RPROGRAM affects:

- The Data Set Report (see “The Data Set Report” on page 82 for more information)
- The GBMAX Report (see “The Maximum Gigabyte Report” on page 88 for more information)
- The Volume Report (see “The Volume Report” on page 101 for more information)
- Estimate Report (see “The Estimate Report” on page 106 for more information)
- IDRC Report (see “The Improved Data Recording Capability Report” on page 120 for more information)
- Management Class Report (see “The Management Class Report” on page 125 for more information)

SPLIT Keyword

Use **SPLIT** to create two data sets that are subsets of the GFTAXTR output data set. One data set contains all the included data sets and the other data set contains all the excluded data sets. Using **SPLIT** maintains the structures of the SMF records, which you cannot do if you sub-divide the SMF input file yourself.

The two split data sets continue to be subsets of the original GFTAXTR output data set. You cannot recombine the subsets.

Abbreviation: None

Syntax

Figure 74 shows the syntax for the **SPLIT** keyword.

[SPLIT]

Figure 74. SPLIT Keyword Syntax

SPLIT has no sub-keywords.

Note: Only use the **SPLIT** keyword once to split the GFTAXTR output data set.

The Importance of SPLIT

SPLIT improves the performance of the volume mount analyzer for those installations that use lengthy include/exclude filter lists during the volume mount analyzer study. Long filter lists cause processing overhead.

Because GFTAVMA compares each filter specified against each SMF 14/15 record in the input file, it is not unusual for some installations to have GFTAVMA runs with excessive elapsed time. Of course this excessive elapsed time is also dependent upon large numbers of SMF 14/15 records that are passed through the filters.

SPLIT

To eliminate the overhead during the GFTAVMA runs, use the **SPLIT** keyword to split the GFTAXTR output data set into two subset data sets. When specified, **SPLIT** directs the SMF 14/15 records into either the include or exclude data set.

After using **SPLIT**, you do not have to worry about the excluded records again or pay for the excessive filtering overhead. The subsequent GFTAVMA runs are against one of the two **SPLIT** output data sets created during the initial **SPLIT** run.

Additional JCL Needed for SPLIT

The include and exclude data sets are identified by two JCL DD statements in the GFTAVMA JCL job stream. Their names must be VMAINCL and VMAEXCL. In the JCL, the GFTAVMA PROC has assigned these DD statements as dummy. In order to use **SPLIT**, you must override these dummy assignments.

To override the current GFTAVMA DD assignments, use the following DD statements to redirect them to a DASD or tape data set. Put the DD statements after the EXEC GFTAVMA statement in the sample JCL.

```
//VMA.VMAINCL DD DSN=your.sequential.data.set.for.includes,
//           DCB=(LRECL=316,RECFM=VB,DSORG=PS),
//           DISP=(NEW,CATLG),
//           SPACE=(.....),
//           UNIT=...
//*
//VMA.VMAEXCL DD DSN=your.sequential.data.set.for.excludes,
//           DCB=(LRECL=316,RECFM=VB,DSORG=PS),
//           DISP=(NEW,CATLG),
//           SPACE=(.....),
//           UNIT=...
```

See the complete JCL in “Submitting the GFTAVMA JCL” on page 48.

Default

SPLIT does not have a default. If you do not specify **SPLIT**, GFTAVMA does not create the two data sets for that run.

Reports Affected

SPLIT affects all GFTAVMA reports.

TAPEDEV Keyword

Use **TAPEDEV** to specify the target tape device to which GFTAVMA should model the data. GFTAVMA uses the target tape device to calculate benefits of IDRC compaction and larger effective device capacity.

Abbreviation: TAPE

Syntax

Figure 75 shows the syntax for the **TAPEDEV** keyword.

`[TAPEDEV (3480 | 348X | 3490E | 3590)]`

Figure 75. **TAPEDEV** Keyword Syntax

The valid tape devices are:

3480

Refers to the base 3480 device without the IDRC feature.

348X

Refers to the 3480 device with the IDRC feature or the base 3490 device.

3490E

Refers specifically to the extended capability 3490E device.

3590

Refers to the Magstar 3590 Tape Drive.

Default

The default for TAPEDEV is 3590.

Relationship to Other Keywords

Use the TAPEDEV keyword with the IDRCFACTOR keyword to project the number of mounts and volumes that could be reduced by using IDRC and the specified tape device.

To specify the double capacity cartridge for 3490E, the Enhanced Capacity Cartridge System Tape, you must also specify the TAPELEN keyword.

Reports Affected

TAPEDEV affects:

- The Estimate Report (see “The Estimate Report” on page 106 for more information)
- The IDRC Report (see “The Improved Data Recording Capability Report” on page 120 for more information)
- The Management Class Report (see “The Management Class Report” on page 125 for more information)

TAPELEN Keyword

Use TAPELEN to specify whether or not GFTAVMA should model to the Enhanced Capacity Cartridge System Tape.

Abbreviation: TLEN

Syntax

Figure 76 shows the syntax for the TAPELEN keyword.

[TAPELEN (1 | 2)]

Figure 76. TAPELEN Keyword Syntax

The valid values for TAPELEN are:

- 1 Refers to a normal tape cartridge with a base capacity of 416 MB for a 3490E tape. This size is the cartridge capacity without IDRC and with 36 tracks. When IDRC is enabled, you can achieve more than the cartridge capacity.
- 2 Refers to the Enhanced Capacity Cartridge System Tape, which is twice the

TAPELEN

length of the tape on a physical cartridge. With the Enhanced Capacity Cartridge System Tape, the effective capacity of a tape cartridge is up to 2500MB.

See *z/OS DFSMS: Implementing System-Managed Storage* for more information on advanced cartridge hardware.

Default

The default for **TAPELEN** is 1.

Relationship to Other Keywords

If you specify **TAPELEN (2)**, you must also specify **TAPEDEV (3490E)**. The **TAPELEN** keyword has no effect on other **TAPEDEV** sub-keywords.

Reports Affected

TAPELEN affects:

- The Estimate Report (see “The Estimate Report” on page 106 for more information)
- The IDRC Report (see “The Improved Data Recording Capability Report” on page 120 for more information)
- The Management Class Report (see “The Management Class Report” on page 125 for more information)

VSORT Keyword

Use **VSORT** to specify whether the list of volume serials in the Volume Report should be sorted by MRV, volume name, or size.

Abbreviation: VSR

Syntax

Figure 77 shows the syntax for the **VSORT** keyword.

`[VSORT (MRV | VOLUME | SIZE)]`

Figure 77. VSORT Keyword Syntax

MRV

Prints a list of volumes in order by MRV. MRV shows the volumes that are mounted frequently but contain small amounts of data.

MRV = number of megabytes divided by the number of mounts.

VOLUME (VOL)

Prints the list of volumes in ascending order by volume serial number. If the volume serial numbers of tapes are meaningful, **VOLUME** is the most useful option.

SIZE (SIZ)

Prints the list of volumes in order by size of the data transferred to or from the volume.

Default

The default for **VSORT** is VOL, which sorts the volume serial names in order by the volume serial number.

Reports Affected

VSORT only affects the Volume Report.

See “The Volume Report” on page 101 for an example of the Volume Report sorted by **VOL**.

VSORT

Appendix A. A Quick Reference to Volume Mount Analyzer Keywords

The following figures show the keyword and sub-keyword abbreviations, the defaults, and the page number where the details of the complete keyword are located.

Abbreviations and Defaults for GFTAXTR Keywords

Table 7 summarizes the GFTAXTR keywords.

Table 7. Quick Reference to GFTAXTR Keywords

Keyword	Abbreviation	Sub-Keyword Abbreviations	Default	Page Number
ACCTFLD	AFLD	(nn)	1	15
DATE	none	(fromdate, todate)	include all	16
LINES	none	(nn)	60	17
TIME	none	(fromtime, totime)	include all	17

Abbreviations and Defaults for Keywords That Filter Your Input

Table 8 summarizes the filtering keywords.

Table 8. Quick Reference for Keywords that Filter Your Input

Keyword	Abbreviation	Sub-Keyword Abbreviations	Default	Page Number
ACCOUNT	ACCT	(INC, EXC)	include all	64
DATASET	DSN	(INC, EXC)	include all	65
DATE	none	(fromdate, todate)	include all	66
EXPDT	EXP	(INC, EXC)	include all	67
FILE	none	(SING MULT BOTH)	BOTH	68
JOBNAME	JOB	(INC, EXC)	include all	69
MAXSIZE	MAX	(nnnn)	infinity	70
MINSIZE	MIN	(nnnn)	zero	70
MOUNT	MOU	(SING MULT BOTH)	BOTH	72
PROGRAM	PGM	(INC, EXC)	include all	73
SYSTEMID	SYID	(INC, EXC)	include all	74
TIME	none	(fromtime, totime)	include all	74
UNIT	none	(INC, EXC)	include all	75
UNITADDR	UADR	(INC, EXC)	include all	76
USAGE	USE	(ALL ACTV BKUP BCOP TEMP SING HSM NONHSM)	ALL	78

Table 8. Quick Reference for Keywords that Filter Your Input (continued)

Keyword	Abbreviation	Sub-Keyword Abbreviations	Default	Page Number
VOLUME	VOL	(INC, EXC)	include all	79

Abbreviations and Defaults for Keywords That Tailor Output Reports

Table 9 summarizes the output report keywords.

Table 9. Quick Reference for Keywords that Tailor Your Output Reports

Keyword	Abbreviation	Sub-keyword Abbreviations	Default	Page Number
BLKSIZE	BLKS	(nnnnn)	32 760	132
BPROGRAM	BPGM	(INC, EXC)	none	133
CHART	none	no sub-keywords	none	134
CLASSIFY	CLASS	ACTV BKUP	none	135
COMPACTION	COMP	(nn)	50%	136
CUTOFF	CUT	(nnnnnnnnnn)	infinity	138
DASDDEV	DASD	(3380S 3380D 3380J 3380E 3380K 3390-1 3390-2 3390-3 3390-9 9345-1 9345-2)	3390-3	138
DSORT	DSR	(DSN JOB PGM SIZ EXP)	DSN	139
FORCECAP	FCAP	no sub-keywords	FORCECAP	140
NOFORCECAP	NOFCAP	no sub-keywords	FORCECAP	140
FSPACE	FSP	(nn)	5%	141
GDG	none	(SHOW GRP)	SHOW	142
IDRCFACTOR	IDRC	(nn.n)	3.0	143
LARGE	LRG	(nnnnnnnnnn)	600MB	144
LEVEL0AGE	L0AGE	(nnnn)	1 day	145
LEVEL1AGE	L1AGE	(nnnn)	1 day	146
LINES	none	(nn)	60 lines	147
L1FSPACE	L1FSP	(nn)	5%	147
REPORT	REP	(DSN EST GB IDRC MC TOP USE VOL)	none	148
RPROGRAM	RPGM	(INC, EXC)	none	150
SPLIT	none	no sub-keywords	none	151
TAPEDEV	TAPE	(3480 348X 3490E 3590)	3490E	152
TAPELEN	TLEN	(1 2)	1	153
VSORT	VSR	(MRV VOL SIZ)	VOL	154

Appendix B. Volume Mount Analyzer Keyword Syntax

The following figures show the syntax for all volume mount analyzer keywords.

Summary of GFTAXTR Keyword Syntax

Table 10 summarizes GFTAXTR keywords and their syntax. All GFTAXTR keywords are optional.

Table 10. Keyword Summary for GFTAXTR

ACCTFLD (*nn*)

DATE (*fromdate[,todate]*)

LINES (*nn*)

TIME (*fromtime[,totime]*)

Summary of GFTAVMA Keyword Syntax

Table 11 summarizes GFTAVMA keywords and their syntax. All GFTAVMA keywords are optional.

Table 11. Keyword Summary for GFTAVMA

ACCOUNT (**INCLUDE** (*account number include list*),
EXCLUDE (*account number exclude list*))

BLKSIZE (*nnnnn*)

BPROGRAM (**INCLUDE** (*backup program include list*),
EXCLUDE (*backup program exclude list*))

CHART

CLASSIFY (**ACTIVE** | **BACKUP**)

COMPACTION (*nn*)

CUTOFF (*nnnnnnnnnn*)

DASDDEV (**3380S** | **3380D** | **3380J** | **3380E** | **3380K**
| **3390-1** | **3390-2** | **3390-3**
| **3390-9** | **9345-1** | **9345-2**)

DATASET (**INCLUDE** (*data set filter list*)
EXCLUDE (*data set filter list*))

DATE (*fromdate[,todate]*)

DSORT (**DATASET** | **JOBNAME** | **PROGRAM** | **SIZE** | **EXPDT**)

EXPDT (**INCLUDE** (*nnnnn*),
EXCLUDE (*yyyyyy*))

FILE (**SINGLE** | **MULTIPLE** | **BOTH**)

FORCECAP | **NOFORCECAP**

FSPACE (*nn*)

GDG (**SHOW** | **GROUP**)

IDRCFACTOR (*nn.n*)

Table 11. Keyword Summary for GFTAVMA (continued)

JOBNAME (INCLUDE (job name filter list), EXCLUDE (job name filter list))
LARGE (nnnnnnnnnn)
LEVEL0AGE (nnnn)
LEVEL1AGE (nnnn)
LINES (nn)
L1FSPACE (nn)
MAXSIZE (nnnn) MINSIZE (nnnn)
MOUNT (SINGLE MULTIPLE <u>BOTH</u>)
PROGRAM (INCLUDE (program filter list), EXCLUDE (program filter list))
REPORT (DATASET,ESTIMATE,GBMAX,IDRC,MGMTCLAS, TOP,USAGE,VOLUME)
RPROGRAM (INCLUDE(rprogram include list), EXCLUDE (rprogram exclude list))
SPLIT
SYSTEMID (INCLUDE (system-id1), EXCLUDE (system-id2))
TAPEDEV (3480 348X 3490E <u>3590</u>)
TAPELEN (<u>1</u> 2)
TIME (fromtime[,totime])
UNIT (INCLUDE (3420,3480,348X,3490E,3590), EXCLUDE (3420,3480,348X,3490E,3590))
UNITADDR (INCLUDE (xxx), EXCLUDE (yyy))
USAGE (<u>ALL</u> ACTIVE, BACKUP,BCOPY,TEMP,SINGLE HSM NONHSM)]
VOLUME (INCLUDE (volume filter list), EXCLUDE (volume filter list))
VSORT (MRV <u>VOLUME</u> SIZE)

Appendix C. Example of Chart Output

Figure 78 on page 163 is an example of the output you would receive if you specify the **CHART** keyword. See “CHART Keyword” on page 134. The output used to create the charts is produced *after* filtering, which allows you to filter against a particular set of data and determine how well that data set is using tape.

You input this sequential data set into your graphics package to create line graphs, bar charts, and pie charts about the data the volume mount analyzer presents in its reports.

The sections of the output are:

CHARTV1

Creates a line graph that shows the cumulative volume and mount percentages by size range (up to 600MB). These numbers come from the Tape Volume Size Range Report part of the GFTASRT2 summary report (See “GFTASRT2—Tape Volume Size Range Report” on page 55). CHARTV1 is produced for every run if you have specified the **CHART** keyword.

CHARTD1

Creates a line graph that shows the tape data set and mount percentages by size range (up to 600MB). These numbers come from the Tape Data Set Size Range Report in the GFTASRT3 summary report (see “GFTASRT3—Tape Data Set Size Range Report” on page 59). CHARTD1 is produced for every run if you have specified the **CHART** keyword.

CHARTD2

Creates a bar chart that shows the single-volume versus multivolume number of data sets, gigabytes, mounts, and volumes. These numbers come from the Tape Data Set Statistics part of the GFTASRT3 summary report (see “GFTASRT3—Tape Data Set Statistics Report” on page 57). CHARTD2 is produced for every run if you have specified the **CHART** keyword.

CHARTD3

Creates a bar chart that shows the single-file versus multifile number of data sets, gigabytes, mounts, and volumes. These numbers come from the Tape Data Set Statistics part of the GFTASRT3 summary report (see “GFTASRT3—Tape Data Set Statistics Report” on page 57). CHARTD3 is produced for every run if you have specified the **CHART** keyword.

CHARTG1

Creates a line graph that shows the maximum hourly gigabytes versus total daily gigabytes for each day. These numbers come from the Maximum Gigabyte Report (see “The Maximum Gigabyte Report” on page 88). CHARTG1 is produced when you specify **CHART** and **REPORT(GBMAX)**.

CHARTG2

Creates a bar chart that shows the maximum hourly gigabytes allocated each hour of the day. These numbers come from the Maximum Gigabyte Report (Allocations by Hour). See “The Maximum Gigabyte Report” on page 88. CHARTG2 is produced when you specify **CHART** and **REPORT(GBMAX)**.

CHARTE1

Creates a bar chart that shows the potential tape mount/volume savings by using 3490E or DFSMS, for example. These numbers come from the Statistical

Mount Savings Breakdown and Statistical Volume Savings Breakdown in the Estimate Report (see “The Estimate Report—Statistical Mount Savings Breakdown” on page 108 and “The Estimate Report—Statistical Volume Savings Breakdown” on page 110). CHARTE1 is produced when you specify **CHART** and **REPORT(EST)**.

CHARTE2

Creates a pie chart that shows the mount percent breakdown by tape mount management category. These numbers come from the Statistical Mount Savings Breakdown part of the Estimate Report (see “The Estimate Report—Statistical Mount Savings Breakdown” on page 108). CHARTE2 is produced when you specify **CHART** and **REPORT(EST)**.

CHARTE3

Creates a pie chart that shows the volume percent breakdown by tape mount management category. These numbers come from the Statistical Volume Savings Breakdown part of the Estimate Report (see “The Estimate Report—Statistical Volume Savings Breakdown” on page 110). CHARTE3 is produced when you specify **CHART** and **REPORT(EST)**.

CHARTE4

Creates a pie chart that shows the data set percent breakdown by tape mount management category. These numbers come from the DASD Buffer Cost Breakdown part of the Estimate Report (see “The Estimate Report—DASD Buffer Cost Breakdown” on page 113). CHARTE4 is produced when you specify **CHART** and **REPORT(EST)**.

CHARTE5

Creates a pie chart that shows the megabyte percent breakdown by tape mount management category. These numbers come from the DASD Buffer Cost Breakdown part of the Estimate Report (see “The Estimate Report—DASD Buffer Cost Breakdown” on page 113). CHARTE5 is produced when you specify **CHART** and **REPORT(EST)**.

```

/ SMF INPUT SAMPLE: 33 DAYS
/      START: 05/27/2003 -- 12:02:01 A.M.
/      END: 06/28/2003 -- 11:59:18 P.M.
/ GFTAXTR RUN:
/      START: 05/21/2004 -- 02:33:02 P.M.
/      END: 05/21/2004 -- 02:57:22 P.M.
/ JES SYSTEM ID(S):
/      SYSTEM ID   OPER SYS
/      -----
/      SYSA       MVS/ESA
/CHARTV1 LINE TPVLSZRG - Cumulative vols/mnts % by size range (up to 600MB)
/TITLE  Tape Vol/Mounts By Size
/YTITLE Cum Percent
/XTITLE Vol Size (MB)
/COLDEF Y1 Y2 X1
/COLNAM  Cum_Vols  Cum_Mnts
*         2.6      2.3      0
*        33.6     34.6      1
*        42.7     43.5      5
*        49.4     50.4     10
*        58.2     59.0     25
*        62.8     63.9     40
*        64.7     65.8     50
*        70.5     71.6     75
*        75.6     76.7    100
*        78.6     79.6    125
*        80.4     81.8    150
*        83.0     84.3    175
*        84.5     85.6    200
*        93.4     93.7    300
*        97.3     97.4    400
*        98.2     98.2    500
*        99.2     99.2    600
/CHARTD1 LINE TPDSSZRG - Tape ds/mnts % by size range (up to 600MB)
/TITLE  Tape Ds/Mounts By Size
/YTITLE Cum Percent
/XTITLE Ds Size (MB)
/COLDEF Y1 Y2 X1
/COLNAM  Cum_Ds  Cum_Mnts
*         6.2      3.3      0
*        58.0     42.5      1
*        69.9     51.8      5
*        75.2     57.3     10
*        82.1     66.1     25
*        86.7     70.9     40
*        88.4     72.7     50
*        91.7     77.4     75
*        93.4     80.0    100
*        94.7     81.8    125
*        95.0     82.3    150
*        95.6     83.2    175
*        96.4     84.1    200
*        97.8     86.1    300
*        98.4     87.7    400
*        98.6     88.2    500
*        99.0     89.6    600

```

Figure 78. Example Chart Output (Part 1 of 4)

```

/CHARTD2 BAR TPDSSVMV - Single-volume vs multi-volume #ds, #GB, #mnts, #vols
/TITLE Single Vs Multi-Volume
/YTITLE Percent
/COLDEF Y1 Y2 LAB
/COLNAM SV MV
* 93.6 6.4 #Ds
* 46.2 53.8 #GB
* 84.4 15.6 #Mnts
* 85.4 14.6 #Vols
/CHARTD3 BAR TPDSSFMF - Single-file vs multi-file #ds, #GB, #mnts, #vols
/TITLE Single Vs Multi-File
/YTITLE Percent
/COLDEF Y1 Y2 LAB
/COLNAM SF MF
* 24.1 75.9 #Ds
* 48.0 52.0 #GB
* 68.0 32.0 #Mnts
* 67.0 33.0 #Vols
/CHARTG1 LINE MAXGBDAY - Max hourly GB versus total daily GB for each day
/TITLE Max Hourly/Total GB
/YTITLE GB
/XTITLE Sample Day
/COLDEF Y1 Y2 LAB
/COLNAM Max_GB Dly_GB
* 6 20 1
* 8 51 2
* 8 62 3
* 9 55 4
* 10 60 5
* 18 125 6
* 9 25 7
* 8 54 8
* 8 57 9
* 8 70 10
* 10 64 11
* 10 67 12
* 18 130 13
* 7 28 14
* 9 54 15
* 9 59 16
* 8 59 17
* 9 57 18
* 9 60 19
* 18 114 20
* 9 42 21
* 9 55 22
* 6 70 23
* 8 68 24
* 9 55 25
* 10 60 26
* 18 120 27
* 9 19 28
* 8 51 29
* 7 67 30
* 10 59 31
* 9 60 32
* 8 64 33

```

Figure 78. Example Chart Output (Part 2 of 4)

```

/CHARTG2 BAR MAXGBHR - Maximum hourly gigabytes allocated each hour of the day
/TITLE Max Hourly GB (By Hour)
/YTITLE Max Hourly GB
/XTITLE Hour
/COLDEF Y1 LAB
/COLNAM Max_GB
*      5  0
*      6  1
*      9  2
*      9  3
*      5  4
*      9  5
*      8  6
*     18  7
*     18  8
*     18  9
*     18 10
*     18 11
*     13 12
*     12 13
*      3 14
*      4 15
*      3 16
*      3 17
*      4 18
*      9 19
*     10 20
*      6 21
*      5 22
*      6 23
/CHART1 BAR TAPESAVE - Potential tape mount/volume savings (3490E , DFSMS)
/TITLE Potential Tape Savings
/COLDEF Y1 Y2 Y3 LAB
/COLNAM Current Direct_3490E DFSMS/3490E
*      38300      36630      7037 Mounts
*      6315      6081      1601 Volumes
/CHART2 PIE TMMNTS - Mount percent breakdown by TMM category
/TITLE TMM Mount Statistics
/COLDEF LAB Y1
/COLNAM Mounts
* Bkup      25476
* Temp      242
* Larg      4050
* Actv      2088
* A_GDG     6443
/CHART3 PIE TMMVOLS - Volume percent breakdown by TMM category
/TITLE TMM Volume Statistics
/COLDEF LAB Y1
/COLNAM Volumes
* Bkup      4987
* Temp      48
* Larg      557
* Actv      86
* A_GDG     635

```

Figure 78. Example Chart Output (Part 3 of 4)

```

/CHART4 PIE TMMDSNS - Data set percent breakdown by TMM category
/TITLE TMM Data Set Statistics
/COLDEF LAB Y1
/COLNAM Data_Sets
* Bkup      38794
* Temp      141
* Larg      418
* Actv      200
* A_GDG     1685
/CHART5 PIE TMMMB - Megabyte percent breakdown by TMM category
/TITLE TMM Megabyte Statistics
/COLDEF LAB Y1
/COLNAM MB
* Bkup      762005
* Temp      3039
* Actv      6779
* A_GDG     71285
* Larg      466686
/ GFTAVMA PROGRAM COMPLETED AT 02:19:32 P.M. ON 09/03/2004 -- RETURN CODE = 0

```

Figure 78. Example Chart Output (Part 4 of 4)

Appendix D. Volume Mount Analyzer Messages

The volume mount analyzer messages appear in GFTAXTR and GFTAVMA output reports.

Table 12. Severity Levels for the Volume Mount Analyzer Messages

Suffix Character	Return Code	Explanation
W	4	<p>A Warning Message indicates either:</p> <ul style="list-style-type: none"> The program is syntactically or semantically incorrect, but the volume mount analyzer can correct the error by extending the statement or removing an option specified on the statement. SMF records were incomplete or contained conflicting information, so the volume mount analyzer dropped them. <p>These are informational messages.</p>
E	8	An Error Message indicates that the volume mount analyzer has detected an incorrect statement and attempted to correct the statement, but was unable to. The program will terminate.
S	12	A Severe Error Message indicates that the volume mount analyzer has detected an incorrect statement, but correcting the statement is impossible. The program will terminate.
T	16	A Terminating Error Message indicates that the volume mount analyzer has detected an error and cannot continue. The program will terminate.

GFTA001E ALL INPUT RECORDS HAVE BEEN EXCLUDED DUE TO FILTERING CONTROLS--NO OUTPUT RECORDS WRITTEN--PROCESSING IS TERMINATED

Explanation: Input was over-filtered.

User Response: Correct the specified filters, and submit the volume mount analyzer job again.

GFTA002T UNABLE TO BUILD THE VOLUME STATISTICS TABLE--NEEDED xxxxxxxx BYTES

Explanation: GFTAVMA needs additional storage space to complete building the in-storage table.

User Response: Increase the REGION size specified on the EXEC JCL statement, using the REGVMA parameter.

GFTA003S NOT ENOUGH VIRTUAL STORAGE FOR THE VOLUME/DATA SET TABLE--NEEDED xxxxxxxx BYTES--PROCESSING TERMINATED

Explanation: GFTAVMA needs additional storage space to complete building the in-storage table.

User Response: Increase the REGION size specified on the EXEC JCL statement.

GFTA004T NO OUTPUT RECORDS WERE WRITTEN TO THE 'XTRCIN' FILE--PROCESSING TERMINATED

Explanation: No input records were selected for processing, thus no output was produced to the GFTAXTR file.

User Response: Verify SMF input contains required record types for tape. Verify that TIME/DATE filtering did not filter out all records. Check for other severe messages.

GFTA005T NO DATA SET OPEN RECORDS (TYPE 14/15) RECORDS WERE WRITTEN TO THE 'XTRCIN' FILE--PROCESSING TERMINATED

Explanation: Some SMF records might have been processed, but there were no OPEN (type 14/15) records for tape.

User Response: Verify SMF input contains required record types for tape. Verify that TIME/DATE filtering did not filter out all records. Check for other severe messages.

GFTA006E NO OUTPUT VOLUME RECORDS WERE WRITTEN TO THE 'XTRCIN' DD--PROCESSING TERMINATED

Explanation: GFTAVMA input processing did not find any volume records in the input file. These records are produced when GFTAXTR successfully pairs the different required record types. Otherwise, the unmatched SMF records are discarded.

User Response: Check the SMF input to ensure that all required SMF record types are present and represent the same time period.

GFTA007T 'XTRIN' FILE CONTAINS INVALID DATA--VERIFY THAT INPUT WAS BUILT BY THE 'GFTAXTR' PROGRAM

Explanation: The input file contains data that is not in the format that GFTAVMA can use.

User Response: Check that XTRIN DD statement points to the correct file.

GFTA008T 'XTRIN' FILE WAS BUILT BY GFTAXTR VERSION 'xxx'--EXPECTED VERSION 'yyy'--PROCESSING TERMINATED

Explanation: The file containing the extracted SMF data is not valid for the current level of GFTAVMA.

User Response: Correct the JCL to point to a file built by the current level of GFTAXTR, or run GFTAXTR again.

GFTA009T 'XTRIN' FILE WAS EMPTY--PROGRAM HIT AN UNEXPECTED END OF FILE--PROCESSING TERMINATED

Explanation: There is no data in the file GFTAXTR built.

User Response: Correct the JCL to point to the correct file containing the extracted data, or run GFTAXTR again.

GFTA010W NO VOLUMES CONTAINING DATA SETS WERE FOUND IN THE INPUT SAMPLE-- GFTASRT2 PROCESSING TERMINATED

Explanation: GFTAVMA was trying to build an input work file for the volume processing phase (GFTASRT2), but no volume records were produced. GFTASRT2 processing terminates because there is nothing to report. Processing continues and goes to GFTASRT3 processing phase.

User Response: Verify that all data sets were not filtered out in the input filtering phase (GFTASRT1) of GFTAVMA.

GFTA011T UNEXPECTED END OF FILE OCCURRED WHILE READING THE 'SORTOUT' FILE--PROCESSING TERMINATED

Explanation: The file to be sorted is empty. This might be an error on the control cards that caused all records to be filtered out.

User Response: Review filters in the control card input. Check for SORT errors (ICE xxx).

GFTA012W JES SYSTEM ID TABLE OVERFLOW (MAX OF 256 SYSTEM IDS ALLOWED)-- ID ='sysid' WAS DROPPED FROM SYSTEM HEADER RECORD

Explanation: There is a "built-in" maximum number of system IDs allowed. The input to this run exceeded this number. The entry is dropped from the header record.

User Response: This is a warning message. If the maximum needs to be increased, this is a new programming request.

GFTA013T EXPECTED 'n' SYSTEM ID(S)--SYSTEM ID RECORD(S) MISSING FROM 'XTRIN' FILE--PROCESSING TERMINATED

Explanation: The list of system IDs to GFTAVMA is kept in a set of header records (records 2 through n), and the number of system IDs is kept in the first record (record type 1). Not all expected system IDs were found.

User Response: Verify that the GFTAXTR run completed successfully and that the GFTAXTR output file has not been modified.

GFTA014E NO DATA SET RECORDS IN THE 'XTRCIN' INPUT FILE FROM 'GFTASRT1'--PROCESSING TERMINATED

Explanation: This message can occur during GFTASRT2 or GFTASRT3 processing. No data records were produced in the previous phase.

User Response: Check for SORT error messages (ICE xxx). Check the input filtering statements to verify that all records were not filtered out.

GFTA015W INVALID RECFM ('yy'X) IN SMF TYPE 'n' RECORD, JOB = 'jobname' DSN = 'dsn'

Explanation: The volume mount analyzer encountered a type 14 or 15 SMF record with an invalid record format value and was unable to correct it. The record is automatically dropped from the sample.

User Response: If this happens often, you could have an SMF problem. Not much can be done about this other than checking the job that created the data set to determine if the error is with the data set or is with SMF, and if necessary, contacting your IBM support center.

GFTA016T ERROR ENCOUNTERED DURING A COMMON FILTER SERVICE (CFS) xxxxxxxxxx CALL, RETURN='yy' REASON='zz'

Explanation: Under normal circumstances, this error should never happen. Most likely it is a program error in the volume mount analyzer or in common filter services. The run forces ABEND 900.

User Response: Contact volume mount analyzer support. Provide dump (SYSUDUMP).

GFTA017W INVALID BLKSIZE VALUE ('blk'X) IN SMF TYPE 'n' RECORD, JOB = 'jobname' DSN = 'dsn'

Explanation: The volume mount analyzer encountered a type 14 or 15 SMF record with an invalid BLOCKSIZE field in the record and was unable to correct it. The record is automatically dropped from the sample.

User Response: If this happens often, you could have an SMF problem. Not much can be done about this other than checking the job that created the data set to determine if the error is with the data set or is with SMF, and if necessary, contacting your IBM support center.

GFTA018T EXPECTED 'n' SYSTEM ID(S)--HIT AN UNEXPECTED END OF FILE ON THE 'XTRIN' FILE--PROCESSING TERMINATED

Explanation: This message can occur during GFTASRT2 or GFTASRT3 processing. A premature EOF occurred before all SYSID records were read in.

User Response: Check for SORT error messages (ICE xxx). Check the input filtering statements to verify that all records were not filtered out.

GFTA020T THERE WAS NOT ENOUGH STORAGE TO PROCESS THE INPUT CONTROL STATEMENTS--NEEDED 'n' BYTES

Explanation: Not enough storage available.

User Response: Increase REGION size using the REGXTR parameter for the GFTAXTR PROC or REGVMA for the GFTAVMA PROC.

GFTA021S STRING ('yyy') IS TOO LONG--MAX LENGTH = 'n'

Explanation: The indicated string exceeds maximum allowable length.

User Response: Check for invalid syntax/spelling.

GFTA022S TOO MANY NESTING LEVELS WERE ENCOUNTERED IN THE ABOVE CONTROL STATEMENTS

Explanation: Error when specifying the input control statements.

User Response: Correct syntax in the input control statement stream.

GFTA023S UNBALANCED PARENS AFTER THE 'yyy' STRING--CHECK FOR MISSING LEFT PAREN

Explanation: Syntax error.

User Response: Correct syntax in the input control statement stream.

GFTA024S MISSING COMMA OR RIGHT PAREN AFTER THE 'yyy' STRING

Explanation: Syntax error.

User Response: Correct syntax in the input control statement stream.

GFTA025E THE 'CLASSIFY' KEYWORD IS NOT ALLOWED WITH THE 'BPGM' KEYWORD

Explanation: These keywords are mutually exclusive.

User Response: Correct syntax in the input control statement stream.

GFTA026E 'xxx' IS AN INVALID VALUE FOR THE 'yyy' KEYWORD--EXPECTED A NUMERIC VALUE OF EITHER 'NN' OR 'NN.N'

Explanation: An invalid value was specified.

User Response: Correct the value specified on the indicated keyword.

GFTA027E 'xxx' IS INVALID LEVEL 1 KEYWORD--PLEASE CHECK FOR PROPER SPELLING

Explanation: Syntax error.

User Response: Check for correct syntax/spelling of the indicated keyword, and submit the volume mount analyzer job again.

GFTA028S 'xxx' IS AN INVALID LEVEL 'n' KEYWORD UNDER THE 'yyy' KEYWORD--PLEASE CHECK FOR PROPER SPELLING

Explanation: Syntax error.

User Response: Check for correct syntax/spelling of the indicated keyword, and submit the volume mount analyzer job again.

GFTA029E MISSING FILTER LIST ON THE 'sss' SUB-KEYWORD OF THE 'xxx' KEYWORD

Explanation: One or more filters were expected.

User Response: Supply the desired filters on the indicated keyword, and submit the volume mount analyzer job again.

GFTA030E THE 'xxx' KEYWORD IS DUPLICATED SPECIFIED

Explanation: You cannot use a keyword multiple times.

User Response: Remove the unnecessary reference to the keyword, and submit the volume mount analyzer job again.

GFTA031E REQUIRED NUMERIC SUB-PARM IS MISSING ON THE xxxxxxxx KEYWORD SPECIFICATION

Explanation: GFTAVMA did not find an expected numeric sub-parameter.

User Response: Supply the required sub-parameter, and submit the volume mount analyzer job again,

GFTA032E KEYWORD 'xxx' HAS AN INVALID VALUE 'vvv'--EXPECTED ALL DECIMAL NUMERICS

Explanation: The indicated value is not valid on this keyword.

User Response: Correct the value 'vvv' to a decimal number, and submit the volume mount analyzer job again.

GFTA033E KEYWORD 'xxx' HAS AN INVALID VALUE 'vvv'--EXPECTED A DECIMAL VALUE FOLLOWED BY 'KB' OR 'MB'

Explanation: Incorrect syntax for 'vvv'.

User Response: Correct the syntax for the indicated keyword and value, and submit the volume mount analyzer job again.

GFTA034S MISSING KEYWORD BETWEEN THE xxxxxxxx AND yyyyyyyyyy DELIMITERS AFTER KEYWORD zzzzzzzz

Explanation: A keyword had delimiters without an expected sub-keyword.

User Response: Correct the syntax in the specified keyword, and submit the volume mount analyzer job again.

GFTA035S UNBALANCED PARENS IN THE CONTROL STATEMENTS--CHECK FOR MISSING RIGHT PAREN

Explanation: Syntax error.

User Response: Correct syntax in input control statements, and submit the volume mount analyzer job again.

GFTA039E THE 'IDRCFACTOR' VALUE 'nn' IS INVALID--IT MUST BE A DECIMAL VALUE BETWEEN 'xx' and 'yy'

Explanation: Invalid value specified.

User Response: Correct the IDRCFACTOR by supplying a decimal value.

GFTA040E THE 'xxxxxx' PERCENT ('yy') IS INVALID--IT MUST BE AN INTEGER BETWEEN 'n' AND 'm'

Explanation: The xxxxxx(yy) PERCENT value is an unsupported value.

User Response: Correct the numerical value, and submit the volume mount analyzer job again.

GFTA041E BOTH 'FORCECAP' AND 'NOFORCECAP' SPECIFIED--ONLY ONE OF THEM IS ALLOWED

Explanation: Specified keywords are mutually exclusive.

User Response: Remove either the FORCECAP or NOFORCECAP keyword, and submit the volume mount analyzer job again.

GFTA042E SYNTAX ERROR ON THE 'xxx' KEYWORD INCLUDE|EXCLUDE LIST--FILTER KEY = 'zzz'

Explanation: The filter specified in the message was not acceptable to common filter services and was rejected.

User Response: Correct the syntax of the specified filter, and submit the volume mount analyzer job again.

GFTA043E NOT ENOUGH STORAGE TO READ IN THE INPUT CONTROL STREAM (EXCEEDED 'n' BYTES)--PROCESSING TERMINATED

Explanation: REGION size is too small.

User Response: Specify a larger REGION size on the EXEC JCL statement, and submit the volume mount analyzer job again.

GFTA044T NOT ENOUGH STORAGE FOR KEYWORD ARRAY (EXCEEDED 'n' BYTES)--LAST KEY='xxx'

Explanation: REGION size is too small.

User Response: Specify a larger REGION size on the JOB or EXEC JCL statement, and submit the volume mount analyzer job again.

GFTA045E ONLY ONE SUB-KEYWORD IS ALLOWED WITH THE 'xxx' KEYWORD--CHOOSE ONE AND RE-RUN

Explanation: Multiple sub-keywords have been specified.

User Response: Correct the keyword syntax, and submit the volume mount analyzer job again.

GFTA046E THE 'GDG(GROUP)' OPTION IS NOT ALLOWED WITH THE 'REPORT(IDRC or EST)' KEYWORD SPECIFICATION

Explanation: Both the IDRC and Estimate Reports process individual data sets. GDG(GROUP) implies a summary grouping of generations, which conflicts with IDRC processing that is used in both reports; therefore the GDG(GROUP) option is not allowed when the IDRC and Estimate Reports are specified.

User Response: Remove the GDG(GROUP) keyword when requesting either the IDRC or Estimate Reports, and submit the volume mount analyzer job again.

GFTA047E THE 'nnn' VALUE 'vv' IS INVALID--IT MUST BE AN INTEGER BETWEEN 'x' AND 'y'

Explanation: Syntax error--a keyword value is outside the allowable value range.

User Response: Change 'vv' to a valid integer value, and submit the volume mount analyzer job again.

GFTA048E BOTH 'CLASSIFY(ACTIVE)' AND 'CLASSIFY(BACKUP)' SPECIFIED--ONLY ONE OPTION ALLOWED

Explanation: Active and backup are mutually exclusive values.

User Response: Specify the CLASSIFY keyword with only one value, active or backup, and submit the volume mount analyzer job again.

GFTA049E THE "FILF#" VALUE (nnnnnnnnn) IS NOT VALID, IT MUST BE A VALUE BETWEEN 100 and 21474836

Explanation: FILF# refers to an external PARM on the EXEC statement. PARM='FILF#(xxxx)' specifies the maximum number of filter keys ("ABC%%XY.*ABC*.*") expected. The default is 10,000. If more than 10,000 filters are required, then specify the requirement on this PARM.

User Response: Count the number of filters in the input keywords. Specify that count (rounded up) on the FILF# parameter of the EXEC JCL statement.

GFTA050E THE 'FILF#' VALUE ABOVE 'vvvv' IS NOT A DECIMAL NUMBER--RUN TERMINATED

Explanation: FILF# refers to an external PARM on the EXEC statement. PARM='FILF#(xxxx)' specifies the maximum number of filter keys ("ABC%%XY.*ABC*.*") expected. The default is 10,000. If more than 10,000 filters are required, then specify the requirement on this PARM.

GFTA051E • GFTA059E

User Response: The FILT#(vvvv) must be a numerical value. Specify the value, and submit the volume mount analyzer job again.

GFTA051E THE 'FILT#' VALUE ABOVE IS 'n' DIGITS--MAXIMUM VALUE LENGTH IS 'x' DIGITS

Explanation: FILT# refers to an external PARM on the EXEC statement. PARM='FILT#(xxxx)' specifies the maximum number of filter keys ("ABC%%XY.*ABC*.**") expected. The default is 10,000. If more than 10,000 filters are required, then specify the requirement on this PARM.

User Response: The FILT#(vvvv) must be a numerical value. Specify the value, and submit the volume mount analyzer job again.

GFTA052E SYNTAX ERROR IN THE ABOVE EXTERNAL PARAMETER--EXPECTED 'FILT#(NNNNNNNNNN)' KEYWORD AND VALUE

Explanation: FILT# refers to an external PARM on the EXEC statement. PARM='FILT#(xxxx)' specifies the maximum number of filter keys ("ABC%%XY.*ABC*.**") expected. The default is 10,000. If more than 10,000 filters are required, then specify the requirement on this PARM.

User Response: Correct the error in the PARM field on the EXEC DD statement. FILT# is the only acceptable value.

GFTA053E EXTERNAL PARAMETER ERROR--INVALID SPECIFICATION IN PARM FIELD

Explanation: FILT# refers to an external PARM on the EXEC statement. PARM='FILT#(xxxx)' specifies the maximum number of filter keys ("ABC%%XY.*ABC*.**") expected. The default is 10,000. If more than 10,000 filters are required, then specify the requirement on this PARM.

User Response: Correct the error in the PARM field on the EXEC DD statement. FILT# is the only acceptable value.

GFTA054E TOO MANY VALUES IN THE LIST WITH THE TIME|DATE KEYWORD--ONLY THE START/STOP VALUES ALLOWED

Explanation: Too many time or date values were specified on the keyword.

User Response: Correct the indicated keyword, and submit the volume mount analyzer job again.

GFTA055E ERROR ON THE SPECIFICATION ON THE 'DATE' KEYWORD--START DATE 'd1' GREATER THAN STOP DATE 'd2'

Explanation: The volume mount analyzer expects two date values that are of a start and end date format.

User Response: The start date must be less than or equal to the end date. Specify the DATE keyword values, and submit the volume mount analyzer job again.

GFTA056E ERROR ON THE SPECIFICATION ON THE 'TIME' KEYWORD--START TIME 'T1' GREATER THAN STOP TIME 'T2'.

Explanation: The volume mount analyzer program expects two time values that are a starting and ending time format.

User Response: The start time must be less than the end time. Specify the TIME keyword, and submit the volume mount analyzer job again.

GFTA057E THE REPORT(TOP(PERCENT(m))) PERCENT VALUE IS INVALID--IT MUST BE AN INTEGER BETWEEN 'xx' and 'yy'

Explanation: The volume mount analyzer program expects an integer value representing a percentage.

User Response: Correct the specified percent value to indicate an integer value, and submit the volume mount analyzer job again.

GFTA058E THE NUMERIC VALUE 'vvvv' ON THE 'xxxxxx' KEYWORD IS TOO HIGH--MAXIMUM VALUE ALLOWED IS 'nnnn'

Explanation: The value specified on the stated keyword is greater than the maximum allowed size.

User Response: Select a new value for the keyword, and submit the volume mount analyzer job again.

GFTA059E A REQUIRED SUB-KEYWORD IS MISSING ON THE 'xxxxxx' KEYWORD SPECIFICATION--MUST BE SPECIFIED

Explanation: Syntax error. The volume mount analyzer expects to find a sub-keyword on the keyword identified.

User Response: Correct the indicated keyword, and submit the volume mount analyzer job again.

GFTA060E A REQUIRED VALUE IS MISSING ON THE 'xxxxxx' KEYWORD SPECIFICATION--MUST BE SPECIFIED

Explanation: Syntax error. The volume mount analyzer expects to find a value on the indicated keyword.

User Response: Correct the indicated keyword, and submit the volume mount analyzer job again.

GFTA061E ONLY ONE VALUE IS ALLOWED WITH THE 'xxxxxx' KEYWORD--SPECIFICATION IS INVALID

Explanation: Syntax error. GFTAVMA only expects one value on the identified keyword. More than one value was found.

User Response: Correct the indicated keyword, and submit the volume mount analyzer job again.

GFTA062E THE VALUE SPECIFICATION 'vv' on the 'xxxxxx' KEYWORD IS TOO LONG--MAX LENGTH = 'n'

Explanation: Syntax error. The value specified on the identified keyword exceeds the length allowed.

User Response: Correct the indicated value, and submit the volume mount analyzer job again.

GFTA063E THE VALUE SPECIFICATION 'vv' on the 'xxxxxx' KEYWORD IS INVALID -- SPECIFICATION NOT ALLOWED

Explanation: Syntax error.

User Response: Correct the indicated value, and submit the volume mount analyzer job again.

GFTA064E ONLY ONE NUMERIC VALUE ALLOWED ON THE 'xxxxxx' KEYWORD SPECIFICATION-- REMOVE THE VALUES NOT WANTED IN THE LIST

Explanation: Syntax error.

User Response: Correct the indicated keyword, and submit the volume mount analyzer job again.

GFTA065E BOTH THE INCLUDE and EXCLUDE KEYWORDS WERE MISSING ON THE 'xxxxxx' SPECIFICATION--THEY ARE REQUIRED

Explanation: Syntax error.

User Response: Specify either an INCLUDE or

EXCLUDE parameter, and submit the volume mount analyzer job again.

GFTA070T SORT PROCESSING ERROR--UNKNOWN RETURN CODE OF 'xx' FROM SORT-- PROCESSING TERMINATED

Explanation: The volume mount analyzer job terminated because SORT has detected an error condition.

User Response: Check the accompanying SORT messages (ICExxx) for more details about the indicated return code.

GFTA071T SORT PROCESSING ERROR--SORT WAS UNSUCCESSFUL DUE TO SOME SEVERE ERROR CONDITION--SEE SORT MESSAGES

Explanation: The volume mount analyzer job terminated because SORT has detected an error condition.

User Response: Check the accompanying SORT messages (ICExxx) for more details about the error condition.

GFTA072T SORT PROCESSING ERROR--'SYSOUT' FILE MISSING FROM THE JCL STREAM-- PROCESSING TERMINATED

Explanation: The volume mount analyzer job terminated because the called SORT program could not find a required SYSOUT file definition.

User Response: Ensure you used the standard JCL that was packaged with GFTAVMA.

GFTA073W INTERNAL PROGRAM WARNING--HSM DATA SET NOT AUTO DETECTED-- DSN='dsn'

Explanation: The data set indicated was generated by DFSMSHsm, but does not match one of the data set names on the volume mount analyzer's list of known DFSMSHsm data sets.

User Response: None, this is an informational message.

GFTA074T POSSIBLE PROGRAM BUG IN 'GFTAXTR' PROGRAM--VOLUME USAGE INDEX OF ZERO WAS IN INPUT RECORD--USER ABEND FORCED

Explanation: Every usage of a volume (such as OPEN or DEMOUNT) should have a unique, non-zero index value so that pairs of OPENs and demounts can be

GFTA075T • GFTA081W

operated on. These indices are assigned by 'GFTAXTR'. A zero value was found and ABEND 902 was issued.

User Response: Contact volume mount analyzer support.

GFTA075T AN IMPOSSIBLE ERROR CONDITION OCCURRED ON THE TRACK CALCULATION FOR THE DATA SET LISTED ABOVE--USER ABEND FORCED

Explanation: The volume mount analyzer was attempting to calculate size of the data set listed above when an impossible internal program condition occurred. The run forces ABEND 903 or 904.

User Response: Contact volume mount analyzer support.

GFTA076W INVALID CHARACTER FOUND IN POSITION 'n' OF THE ACCOUNTING FIELD--ACCOUNT DATA DROPPED. LENGTH='n' BYTES--ACCOUNT FIELD='yyyyyy'

Explanation: The indicated field of the accounting record contains invalid data. The record has been discarded.

User Response: None, this is an informational message.

GFTA077W SMF INPUT ERROR--THE FOLLOWING DEMOUNT (TYPE 21) RECORD WAS MISSING THE VOLUME INFORMATION--RECORD WAS DROPPED

Explanation: The message is followed by specific information about the dropped record.

User Response: None, this is an informational message.

GFTA078W INVALID SMF RECORD MOVE TIME/DATE STAMP IN A TYPE 'n' RECORD--INVALID FIELD='xxxxxx'. LAST VALID SMF MOVE DATE/TIME WAS 'MM/DD/YYYY--hh:mm:ss'--RECORD DROPPED

Explanation: GFTAXTR detected an invalid time or date value in the SMF record type described in the message. The time and date values are in the standard header portion of the SMF record. GFTAXTR expects all the SMF records to have valid time and date values in these fields. If all the SMF records in the input sample are thrown away, then the input SMF records are probably not in the standard SMF record format. Otherwise, an occasional invalid time or date value could occur in thousands of SMF accounting records.

User Response: If all the SMF records in your input sample are thrown away, check the SMF input data to see if it conforms to the standard SMF formats. If the error is one in thousands, then this message is informational only.

GFTA079W DSN SYNTAX ERROR IN SMF TYPE 'n' RECORD--INVALID CHARACTERS FOUND AT THE END OF THE DATA SET NAME--CHARACTERS IGNORED-- DSN='dsn'

Explanation: The data set name indicated in the message contains extra characters that do not conform with acceptable data set name standards.

User Response: None, this is an informational message.

GFTA080W DSN SYNTAX ERROR IN SMF TYPE 'n' RECORD--RECORD DROPPED, JOB='jobname', DSN='dsn'--

Explanation: An invalid data set name was encountered.

This message is followed by one of four possible sub-messages:

- QUALIFIER NUMBER 'n' IN THE DATA SET NAME ABOVE IS MISSING--- DATA SET QUALIFIERS MUST BE 1 TO 8 BYTES.
- QUALIFIER NUMBER 'n' IN THE DATA SET NAME ABOVE IS TOO LONG-- THE MAXIMUM LENGTH ALLOWED IS 8 BYTES.
- QUALIFIER NUMBER 'n' IN THE DATA SET NAME ABOVE BEGINS WITH A NON-ALPHABETIC CHARACTER--THIS IS INVALID.
- AN INVALID CHARACTER WAS ENCOUNTERED IN POSITION 'n' IN THE DATA SET NAME ABOVE--IT MUST BE ALPHAMERIC.

User Response: None, this is an informational message.

GFTA081W UNKNOWN TAPE UNIT UCCTYPE='uu'X IN SMF TYPE 'n' RECORD, JOB = 'jobname' VOL = 'volser'--RECORD DROPPED

Explanation: The UCB field in the SMF record indicates a unit type that is unrecognizable to GFTAVMA.

User Response: None, this is an informational message.

GFTA082T GETMAIN FAILURE FOR FILTER KEY
LISTS--NEEDED 'n' BYTES--
PROCESSING TERMINATED

Explanation: REGION size is too small.

User Response: Specify a larger REGION size value on the JOB or EXEC JCL statement, and submit the volume mount analyzer job again.

GFTA083T AN I/O ERROR OCCURRED WHILE
ATTEMPTING TO READ THE
'VMAFLTRS' DATA SET. (SEE
FOLLOWING ERROR MESSAGE)

Explanation: All INCLUDE or EXCLUDE filter keys are parsed to an internal format and then written to the data set referenced by the VMAFLTRS DD statement. An I/O error occurred on the file when the volume mount analyzer was re-reading the VMAFLTRS data set. Processing is terminated. There is an accompanying GFTA2xx message with additional information.

User Response: This is usually a temporary data set. Look in the JCL to determine which volume was selected. Check explanation for the GFTA2xx message. Contact the system programmer to determine if these are hardware I/O errors. Contact local IBM support if the problem is with IBM hardware or software.

GFTA084T A GETMAIN FAILURE FOR FILTER
WORK AREA--NEEDED 'n'
BYTES--PROCESSING TERMINATED

Explanation: REGION size is too small.

User Response: Specify a larger REGION size value on the JOB or EXEC JCL statement, and submit the volume mount analyzer job again.

GFTA085S INVALID SYNTAX FOR FILTER KEY:
'ffff'

Explanation: Common filter services found a filter key that was not acceptable.

User Response: Correct the syntax/spelling for the specified filter, and submit the volume mount analyzer job again.

GFTA086W WARNING—ERRORS DETECTED
DURING SIZE CALCULATION ON
THE FOLLOWING DATA SET(S):

Explanation: The volume mount analyzer was trying to calculate the size of the data set in tracks (based on the DASDDEV keyword and the data set attributes, such as LRECL, BLKSIZE, and RECFM). If the volume mount analyzer is unable to perform the calculation because of invalid or missing information, it lists the data set name and all pertinent information in the following format:

```
RC RECFM LRECL BLKSIZE DEVICE #BLOCKS DATA SET NAME
-----
XX XXXX XXXX XXXX XXXX XXXXX XXXX.XXX.XXXX
.
.
.
```

The following values are possible for RC:

- 04** Invalid record format (RECFM) value.
- 08** Impossible condition (should never happen). The volume mount analyzer forces ABEND 903.
- 12** Impossible condition (should never happen). The volume mount analyzer forces ABEND 904.
- 16** Specified LRECL value was zero. Unable to calculate size without LRECL.
- 20** Specified LRECL value is higher than the maximum record size for the target device.
- 24** BLKSIZE value is zero. Unable to calculate size without BLKSIZE.
- 28** BLKSIZE value is greater than the maximum record size for target device.

For RC=8 or RC=12, message GFTA075 is also generated, and the program abends. For all other RC values, the size of the data set is calculated as zero, and the job continues.

User Response: Either ignore, or filter these records out. If necessary, contact local IBM support to determine why the SMF data is in error.

GFTA087E THE 'TIME' VALUE SPECIFIED 'T1' IS
TOO LONG--EXPECTED 'HH:MM:SS'

Explanation: Refer to the TIME keyword description.

User Response: Correct the format of the time specified, and submit the volume mount analyzer job again.

GFTA088E THE 'n' SPECIFICATION IN THE
'TIME' VALUE ('xx') IS INVALID

Explanation: Either the hour, minute or second value supplied is invalid.

User Response: Correct the specified value, and submit the volume mount analyzer job again.

GFTA089T A GETMAIN FAILURE FOR PARSED
KEY AREA--NEEDED '*n*'
BYTES--PROCESSING TERMINATED

Explanation: REGION size is too small.

User Response: Supply a larger REGION size on the JOB or EXEC JCL statement, and submit the volume mount analyzer job again.

GFTA090E INVALID VALUE '*vv*' ON THE 'TIME'
KEYWORD--EXPECTED '*HH:MM:SS*'
'*HH:MM*' OR '*HH*' FORMAT

Explanation: Refer to the TIME keyword description.

User Response: Supply a TIME value that matches the acceptable format, and submit the volume mount analyzer job again.

GFTA091E INVALID VALUE '*vv*' ON THE 'DATE'
KEYWORD--EXPECTED '*MM/DD/YY*',
'*MM/DD/YYYY*', '*YYDDD*', OR
'*YYYYDDD*' FORMAT

Explanation: Refer to the DATE keyword description.

User Response: Supply a DATE value that matches the acceptable format, and submit the volume mount analyzer job again.

GFTA092E INVALID JULIAN 'DAY'
SPECIFICATION '*ddd*' IN THE VALUE
'*xxx*' ON THE 'DATE' KEYWORD

Explanation: The value specified '*ddd*' is not a valid Julian day.

User Response: Correct the specified value.

GFTA093E INVALID JULIAN 'YEAR'
SPECIFICATION '*yy*' IN THE VALUE
'*xxx*' ON THE 'DATE' KEYWORD

Explanation: The value specified '*yy*' is not a valid Julian year.

User Response: Correct the specified value.

GFTA094E INPUT 'DATE' STRING '*d1*' IS TOO
LONG--MAXIMUM LENGTH
ALLOWED IS '*n*'

Explanation: Syntax error.

User Response: Correct the format of the date specified.

GFTA095E INVALID '*n*' SPECIFICATION IN THE
VALUE '*xx*' ON THE 'DATE'
KEYWORD

Explanation: Either the year, month or day has been specified incorrectly.

User Response: Correct the indicated value.

GFTA096W INVALID FORMAT FOR FILTER DATA:
'*ffff*'--RECORD EXCLUDED FROM
INPUT

Explanation: The input SMF fields (such as DSN and EXPDT) might contain unsupported data when passed to common filter services. If so, this warning message is printed and the data is excluded from the run.

User Response: Report the SMF data problem to the local IBM support center.

GFTA097W WRONG LEVEL OF COMMON FILTER
SERVICES (CFS) ENCOUNTERED--
WILL USE INTERNAL FILTERING
ROUTINE (NO PERCENT SIGN
SUPPORT)

Explanation: Probable volume mount analyzer program error. The volume mount analyzer requires that DFP 2.2 or greater be installed. This should have the proper level of CFS. If not, then the volume mount analyzer prints this message and automatically switches to its internal filtering routines to allow the run to continue.

User Response: Determine the DFP level, and contact volume mount analyzer support.

GFTA098T START TIME/DATE OF TAPE
ALLOCATION > END
TIME/DATE--POSSIBLE SMF INPUT
ERROR

Explanation: An invalid condition occurred in a demount record (type 21) in the input file from GFTAXTR--possible program error. The last record written to XTRCIN contains the invalid record. GFTAVMA prints this message and the GFTA099T message and abends with ABEND 901.

User Response: Contact volume mount analyzer support, and provide dump (SYSUDUMP).

GFTA099T "IMPOSSIBLE" PROGRAM
CONDITION OCCURRED--SEE LAST
RECORD IN 'XTRCIN' FILE

Explanation: An invalid condition occurred in a demount record (type 21) in the input file from GFTAXTR--possible program error. The last record written to XTRCIN contains the invalid record. GFTAXTR also prints the GFTA098T message, and abends with ABEND 901.

User Response: Contact volume mount analyzer support, and provide dump (SYSUDUMP).

GFTA101W INTERNAL 'XTRIN' FILE ERROR--VOL='volser', DSN='dsn', OPI#='xx', OPO#='yy',--MOUNTS SET FROM 0 TO 1

Explanation: GFTASRT2 encountered an external SMF record with zero mounts associated. The mount count was set from 0 to 1 and processing continues.

User Response: None, this is an informational message.

GFTA102W SMF INPUT ERROR--THE FOLLOWING TYPE 14/15 RECORD HAS A DATA SET SIZE THAT EXCEEDS THE SIZE OF A GENERAL PURPOSE REGISTER--THE SIZE HAS BEEN RESET TO 2,147,483,647 KILOBYTES--DSN ='dsn'--VOLSER='volser'--TSMV='YYDDD'--'hh:mm:ss'

Explanation: GFTASRT1 encountered an SMF record with a data set size that was unreasonably large. The size was reset to the maximum allowable size.

User Response: None, this is an informational message.

GFTA103W INVALID LENGTH IN ACCOUNT FIELD 'xx' OF THE TYPE 30 RECORDS-- RECORDS CORRECTED BY GFTAXTR

Explanation: The indicated field in the SMF Type 30 record, contained a data field that was not an acceptable length.

User Response: None, this is an informational message.

GFTA104W SMF INPUT ERROR--DEMOUNT RECORD HAS NO MATCHING 14/15 RECORD--THE FOLLOWING DEMOUNT (TYPE 21) RECORD WAS DROPPED:

Explanation: Missing record in the SMF input data. GFTAXTR does not have sufficient information to be able to perform further analysis for this demount record. The record is being discarded. The message is followed by key information fields from the record for identification purposes.

User Response: None, unless there are an unacceptable number of these messages. In that case, ensure that you are collecting the appropriate SMF records, and run the IFASMFDP job again.

GFTA105W SMF INPUT ERROR--MISSING DEMOUNT RECORD--THE FOLLOWING DEMOUNT (TYPE 21) RECORD WAS GENERATED:

Explanation: Missing record in the SMF input data. GFTAXTR has sufficient information from the OPEN record associated with the mount to be able to create a substitute demount record. The message is followed by key information fields from the record for identification purposes.

User Response: None, unless there are an unacceptable number of these messages. In that case, ensure that you are collecting the appropriate SMF records, and run the IFASMFDP job again.

GFTA106W SMF INPUT ERROR ON TYPE 30 RECORD (SUBTYPE 'xx') INVALID OFFSET FIELD TO SUBSYSTEM SECTION--RECORD WAS DROPPED: 1ST 50 BYTES HEX: 'xxxxxxx...'

Explanation: An SMF type 30 record has been found to have an error that cannot be corrected.

User Response: None, this is an informational message.

GFTA107W SMF INPUT ERROR ON TYPE 'n' RECORD INVALID RDW LENGTH

Explanation: The record descriptor work in the indicated record is not of an acceptable length.

User Response: None, this is an informational message.

GFTA108S NOT ENOUGH VIRTUAL STORAGE FOR THE RECALL/TRACKS ANALYSIS TABLE--NEEDED 'nnnnnn' BYTES--PROCESSING TERMINATED

Explanation: GFTASRT3 was in the process of building an internal table to calculate the number of recall tracks necessary. However, a GETMAIN for storage was not successful. The number of bytes required is contained in the "nnnnnn" field of the message.

User Response: Increase the REGION size for the GFTAVMA program, and submit the volume mount analyzer job again.

GFTA109S 'GFTAXTR' INPUT ERROR--THE LAST VERSION OF A DELETED DATA SET HAD AN "OLD" VOLUME COUNT OF ZERO--PROCESSING WILL TERMINATE WITH AN
abend--DSN='dsn', VOLSER=vvvvvv,-- TSMV=yy/mm/dd,hh:mm:ss, -- VSQ#=#nn, DSQ#=#xx, VUIX=yy

Explanation: An internal program error has occurred. The particular field should never have a value of zero. This message is followed by ABEND 906.

User Response: Contact volume mount analyzer support.

GFTA110W UNABLE TO OPEN THE 'VMACHART' FILE TO BUILD THE SEQUENTIAL DATA SET FOR GRAPHIC CHARTS--NO FILE BUILT BUT PROCESSING CONTINUES

Explanation: The CHART keyword was specified in the input stream, but GFTAVMA was unsuccessful when it attempted to open the VMACHART DCB.

User Response: Check the GFTAVMAP JCL for an incorrectly specified VMACHART file. Correct the error, and submit the volume mount analyzer job again.

GFTA111W UNABLE TO OPEN THE *dddddd* FILE TO BUILD THE INCLUDE/EXCLUDE SPLIT FILES--NO FILES BUILT BUT PROCESSING CONTINUES

Explanation: The SPLIT keyword was specified in the input stream, but GFTAVMA was unsuccessful when it attempted to open either the VMAINCL or VMAEXCL DCBs. The JCL DD statement name is identified in the *dddddd* field of the message. The SPLIT keyword requires the definition of these two data sets and JCL DD statements in the GFTAVMAP JCL.

User Response: Check the GFTAVMAP JCL for an incorrectly specified VMAINCL or VMAEXCL DD statement. Correct the error, then submit the volume mount analyzer job again.

GFTA200E UNABLE TO RETRIEVE INFORMATION FROM THE *dddddd* DD SINCE IT WAS MISSING FROM THE JCL

Explanation: The volume mount analyzer program, GFTAXTR or GFTAVMA, attempted to open the data set identified by the *dddddd* name in the message. The JCL DD statement identified by this name must be in the volume mount analyzer JCL, otherwise processing terminates.

User Response: Check the GFTAXTRP or GFTAVMAP JCL for an incorrect or missing JCL DD statement identified by the *dddddd* name. Correct the JCL

problem, and submit the volume mount analyzer job again.

GFTA201E BLKSIZE CALCULATION ERROR--REASON=*rc* (DEC), DD=*dddddd*, RECFM=*recfn*, LRECL=*lrecl*, TARGET BLKSIZE=*blksize*-- PROCESSING TERMINATED

Explanation: The volume mount analyzer program, GFTAXTR or GFTAVMA attempted to assign a BLKSIZE to the data set identified by the *dddddd* JCL DD statement and was not successful. There is a reason code in the message indicating why the BLKSIZE calculation failed. The reason codes are:

- 04 RECFM is invalid
- 08 not a DASD device
- 12 invalid DASD type
- 16 LRECL equals zero
- 20 LRECL is too large

User Response: Correct the error indicated in the reason code on the JCL DD statement identified by the *dddddd* name, and submit the volume mount analyzer job again.

GFTA202E UNABLE TO CALCULATE THE BLKSIZE FOR THE *dddddd* DD SINCE THE DATA SET WAS ALREADY OPEN

Explanation: The volume mount analyzer program, GFTAXTR or GFTAVMA attempted to open the data set identified by the *dddddd* name in the message. The data set identified by the *dddddd* JCL DD statement was currently open. The volume mount analyzer processing terminates.

User Response: This should never occur. Check the GFTAXTRP or GFTAVMAP JCL DD statement identified by the *dddddd* name. The data set named by this DD statement could also be named in another DD statement in the JCL job. Correct the JCL data set naming problem, and submit the volume mount analyzer job again.

GFTA203T INPUT SYNAD I/O ERROR *error text* ON UADR=*unit address*, DD=*ddname*, DSN=*data set name*

Explanation: The volume mount analyzer, GFTAXTR or GFTAVMA encountered an I/O error during QSAM GETL processing. The associated SYNAD routine returned with the error information that is described in the *error text* portion of the message. Additional information such as data set name, JCL DDNAME, and unit address are also contained in the message.

User Response: Probable hardware media error. Move the data set to a different volume, and submit the

volume mount analyzer job again.

GFTA204T **OUTPUT SYNAD I/O ERROR** *error text*
ON UADR=unit address, DD=ddname,
DSN=data set name

Explanation: The volume mount analyzer program, GFTAXTR or GFTAVMA encountered an I/O error during QSAM PUTM processing. The associated SYNAD routine returned with the error information that is described in the "error text" portion of the message. Additional information, such as data set name, JCL ddname, and unit address are contained in the message.

User Response: Probable hardware media error. Move the data set to a different volume, and submit the volume mount analyzer job again.

GFTA205E **UNABLE TO OPEN FILE FOR INPUT**
PROCESSING-- DDNAME=dddddd

Explanation: The volume mount analyzer, GFTAXTR or GFTAVMA attempted to open a data set for input processing, but was not successful. The data set is identified by the *dddddd* name from the JCL DD statement.

User Response: Check the JCL statement identified in the message, and correct the data set problem.

GFTA206E **UNABLE TO OPEN FILE FOR OUTPUT**
PROCESSING-- DDNAME=dddddd

Explanation: The volume mount analyzer program, GFTAXTR or GFTAVMA, attempted to open a data set for output processing, but was not successful. The data set is identified by the *dddddd* name from the JCL DD statement.

User Response: Check the JCL statement identified in the message, and correct the data set problem.

ABEND 900

Explanation: Impossible return code from common filter services.

User Response: See accompanying GFTAxxx error message. Contact volume mount analyzer support.

ABEND 901

Explanation: Tape start time is greater than tape end time.

User Response: See accompanying GFTAxxx error message. Contact volume mount analyzer support.

ABEND 902

Explanation: Vol/Usage index should never equal 0.

User Response: See accompanying GFTAxxx error message. Contact volume mount analyzer support.

ABEND 903

Explanation: The target device code (shown as 'xx'X under 'DEVICE' in previous message) is unsupported. It should always be '20'X for DASD devices.

User Response: See accompanying GFTAxxx error message. Contact volume mount analyzer support.

ABEND 904

Explanation: The DASD device type (shown as 'xx'X under 'DEVICE' in previous message) is not supported.

User Response: See accompanying GFTAxxx error message. Contact volume mount analyzer support.

ABEND 906

Explanation: The "old" volume count has a value of zero. This is an impossible condition.

User Response: See accompanying GFTAxxx error message. Contact volume mount analyzer support.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

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

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States, or other countries, or both:

DB2	IBM
DFSMSdfp	MVS/ESA
DFSMSdss	RACF
DFSMSHsm	TotalStorage
DFSMSrmm	z/OS

Other company, product, and service names may be trademarks or service marks of others.

Glossary

This glossary defines technical terms and abbreviations used in DFSMS documentation. If you do not find the term you are looking for, refer to the index of the appropriate DFSMS manual or view the *Glossary of Computing Terms* located at:

<http://www.ibm.com/ibm/terminology/>

This glossary includes terms and definitions from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42nd Street, New York, New York 10036. Definitions are identified by the symbol (A) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published part of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- The *IBM Dictionary of Computing*, New York: McGraw-Hill, 1994.

The following cross-reference is used in this glossary:

See: This refers the reader to (a) a related term, (b) a term that is the expanded form of an abbreviation or acronym, or (c) a synonym or more preferred term.

A

ACS. See *automatic class selection*

active data. For tape mount management, application data that is frequently referenced, usually small in size,

and managed better in the DFSMS storage hierarchy than on tape. Contrast with *inactive data*.

allocation. Generically, the entire process of obtaining a volume and unit of external storage, and setting aside space on that storage for a data set.

automatic class selection (ACS). A mechanism for assigning SMS classes and storage groups to data sets.

automated tape library data server. A device consisting of robotic components, cartridge storage areas, tape subsystems, and controlling hardware and software, together with the set of tape volumes that reside in the library and can be mounted on the library tape drives. Contrast with *manual tape library*. See also *tape library*.

automatic class selection (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.

automatic dump. In DFSMSHsm, the process of using DFSMSdss automatically to do a full-volume dump of all allocated space on a primary storage volume to designated tape dump volumes.

availability. For a storage subsystem, the degree to which a data set or object can be accessed when requested by a user.

B

backup. The process of creating a copy of a data set or object to be used in case of accidental loss.

base configuration. The part of an SMS configuration that contains general storage management attributes, such as the default management class, default unit, and default device geometry. It also identifies the systems or system groups that an SMS configuration manages.

C

Cartridge System Tape. The base tape cartridge media used with 3480 or 3490 Magnetic Tape Subsystems. Contrast with *Enhanced Capacity Cartridge System Tape*.

class transition. An event that brings about change to an object's service-level criteria, causing OAM to invoke ACS routines to assign a new storage class or management class to the object.

construct. One of the following: data class, storage class, management class, storage group, aggregate group, base configuration.

D

DASD. *see direct access storage device*

DASD volume. A DASD space identified by a common label and accessed by a set of related addresses. See also *volume, primary storage, migration level 1, migration level 2.*

data class. A collection of allocation and space attributes, defined by the storage administrator, that are used to create a data set.

Data Facility Sort. An IBM licensed program that is a high-speed data processing utility. DFSORT provides an efficient and flexible way to handle sorting, merging, and copying operations, as well as providing versatile data manipulation at the record, field, and bit level.

data set. In DFSMS, the major unit of data storage and retrieval, consisting of a collection of data in one of several prescribed arrangements and described by control information to which the system has access. In z/OS non-UNIX environments, the terms *data set* and *file* are generally equivalent and sometimes are used interchangeably. See also *file*. In z/OS UNIX environments, the terms *data set* and *file* have quite distinct meanings.

default unit. Part of the SMS base configuration, it identifies an esoteric (such as SYSDA) or generic (such as 3390) device name. If a user omits the UNIT parameter on the JCL or the dynamic allocation equivalent, SMS applies the default unit if the data set has a disposition of MOD or NEW and is *not* system-managed.

Device Support Facilities (ICKDSF). A program used for initialization of DASD volumes and track recovery.

DFSMSdfp. A DFSMS functional component or base element of z/OS, that provides functions for storage management, data management, program management, device management, and distributed data access.

DFSMSdss. A DFSMS functional component or base element of z/OS, used to copy, move, dump, and restore data sets and volumes.

DFSMSHsm. A DFSMS functional component or base element of z/OS, used for backing up and recovering data, and managing space on volumes in the storage hierarchy.

DFSMSHsm-managed volume. (1) A primary storage volume, which is defined to DFSMSHsm but which does not belong to a storage group. (2) A volume in a storage group, which is using DFSMSHsm automatic

dump, migration, or backup services. Contrast with *system-managed volume, DFSMSrmm-managed volume.*

DFSMSrmm. A DFSMS functional component or base element of z/OS, that manages removable media.

DFSMSrmm-managed volume. A tape volume that is defined to DFSMSrmm. Contrast with *system-managed volume, DFSMSHsm-managed volume.*

DFSORT. *see Data Facility Sort*

direct access storage device (DASD). A mass storage medium on which a computer stores data.

dummy storage group. A type of storage group that contains the serial numbers of volumes no longer connected to a system. Dummy storage groups allow existing JCL to function without having to be changed. See also *storage group.*

E

Enhanced Capacity Cartridge System Tape. Cartridge system tape with increased capacity that can only be used with 3490E Magnetic Tape Subsystems. Contrast with *Cartridge System Tape.*

erase-on-scratch. The physical erasure of data on a DASD data set when the data set is deleted (scratched).

expiration. (1) The process by which data sets or objects are identified for deletion because their expiration date or retention period has passed. On DASD, data sets and objects are deleted. On tape, when all data sets have reached their expiration date, the tape volume is available for reuse. (2) In DFSMSrmm, all volumes have an expiration date or retention period set for them either by vital record specification policy, by user-specified JCL when writing a data set to the volume, or by an installation default. When a volume reaches its expiration date or retention period, it becomes eligible for release.

F

file. A collection of information treated as a unit. In z/OS non-UNIX environments, the terms *data set* and *file* are generally equivalent and are sometimes used interchangeably. See also *data set.*

filtering. The process of selecting data sets based on specified criteria. These criteria consist of fully or partially-qualified data set names or of certain data set characteristics.

G

GB. *See gigabyte*

GDG. *See generation data group*

GDS. Generation data set

generation data group (GDG). A collection of historically related non-VSAM data sets that are arranged in chronological order; each data set is called a generation data set.

gigabyte (GB). One billion (10⁹) bytes.

group. (1) With respect to partitioned data sets, a member and the member's aliases that exist in a PDS or PDSE, or in an unloaded PDSE. (2) A collection of users who can share access authorities for protected resources.

I

ICKDSF. *see Device Support Facilities*

ICL. Integrated cartridge loader

IDCAMS. Access method services

IDRC. *see improved data recording capability*

improved data recording capability (IDRC). A recording mode that can increase the effective cartridge data capacity and the effective data rate when enabled and used. IDRC is always enabled on the 3490E Magnetic Tape Subsystems.

inactive data. (1) A copy of active data, such as vital records or a backup copy of a data set. Inactive data is never changed, but can be deleted or superseded by another copy. (2) In tape mount management, data that is written once and never used again. The majority of this data is point-in-time backups. (3) Objects infrequently accessed by users and eligible to be moved to the optical library or shelf. (4) Contrast with *active data*.

interval migration. In DFSMSHsm, automatic migration that occurs when a threshold level of occupancy is reached or exceeded on a DFSMSHsm-managed volume, during a specified time interval. Data sets are moved from the volume, largest eligible data set first, until the low threshold of occupancy is reached.

J

JCL. *see job control language*

job control language (JCL). A control language used to identify a job to an operating system and to describe the job's requirements.

K

KSDS. *see key-sequenced data set*

key-sequenced data set (KSDS). A VSAM data set whose records are loaded in key sequence and controlled by an index.

M

management class. A collection of management attributes, defined by the storage administrator, used to control the release of allocated but unused space; to control the retention, migration, and backup of data sets; to control the retention and backup of aggregate groups, and to control the retention, backup, and class transition of objects.

MB. *see megabyte*

megabyte (MB). 1) For processor storage, real and virtual storage, and channel volume, 2²⁰ or 1 048 576 bytes. 2) For disk storage capacity and communications volume, 1 000 000 bytes.

migration. The process of moving unused data to lower cost storage in order to make space for high-availability data. If you wish to use the data set, it must be recalled. See also *migration level 1*, *migration level 2*.

migration control data set (MCDS). In DFSMSHsm, a VSAM key-sequenced data set that contains statistics records, control records, user records, records for data sets that have migrated, and records for volumes under migration control of DFSMSHsm.

migration level 1. DFSMSHsm-owned DASD volumes that contain data sets migrated from primary storage volumes. The data can be compressed. See also *storage hierarchy*. Contrast with *primary storage*, *migration level 2*.

migration level 2. DFSMSHsm-owned tape or DASD volumes that contain data sets migrated from primary storage volumes or from migration level 1 volumes. The data can be compressed. See also *storage hierarchy*. Contrast with *primary storage*, *migration level 1*.

modeling. The process of projecting future storage requirements for DASD and tape if data is managed by DFSMS.

Multiple Virtual Storage (MVS). The short name for MVS/XA, MVS/ESA, or the MVS element of the z/OS operating system.

MVS. *See Multiple Virtual Storage*

MVS/ESA. Multiple Virtual Storage/Enterprise Systems Architecture. A z/OS operating system environment that supports ESA/390.

MVS/ESA SP. An IBM licensed program used to control the z/OS operating system. MVS/ESA SP together with DFSMS compose the base MVS/ESA operating environment. See also z/OS.

O

OAM. See *object access method*

object access method (OAM). An access method that provides storage, retrieval, and storage hierarchy management for objects and provides storage and retrieval management for tape volumes contained in system-managed libraries.

object backup storage group. A type of storage group that contains optical or tape volumes used for backup copies of objects. See also *storage group*.

object storage group. A type of storage group that contains objects on DASD, tape, or optical volumes. See also *storage group*.

object storage hierarchy. A hierarchy consisting of objects stored in DB2 table spaces on DASD, on optical or tape volumes that reside in a library, and on optical or tape volumes that reside on a shelf. See also *storage hierarchy*.

optical volume. Storage space on an optical disk, identified by a volume label. See also *volume*.

OS/390. OS/390 is a network computing-ready, integrated operating system consisting of more than 50 base elements and integrated optional features delivered as a configured, tested system. See also *MVS/ESA SP*.

P

partitioned data set (PDS). A data set on direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

partitioned data set extended (PDSE). A data set that contains an indexed directory and members that are similar to the directory and members of partitioned data sets. A PDSE can be used instead of a partitioned data set.

PDS. See *partitioned data set*

PDSE. See *partitioned data set extended*

performance. (1) A measurement of the amount of work a product can produce with a given amount of resources. (2) In a DFSMS environment, a measurement of effective data processing speed with respect to objectives set by the storage administrator. Performance is largely determined by throughput, response time, and system availability.

permanent data set. A user-named data set that is normally retained for longer than the duration of a job or interactive session. Contrast with *temporary data set*.

pool storage group. A type of storage group that contains system-managed DASD volumes. Pool storage groups allow groups of volumes to be managed as a single entity. See also *storage group*.

primary storage. A DASD volume available to users for data allocation. The volumes in primary storage are called primary volumes. See also *storage hierarchy*. Contrast with *migration level 1*, *migration level 2*.

program management. The task of preparing programs for execution, storing the programs, load modules, or program objects in program libraries, and executing them on the operating system.

R

RACF. See *Resource Access Control Facility*

recall mount. A mount that results from a data set being referenced after it has migrated to migration level 1 DASD or migration level 2 tape. There are two types of recall mounts: level 2 to primary DASD, which requires a mount, and level 1 DASD to primary DASD, which only requires data transfer and thus is less costly.

Resource Access Control Facility (RACF). An IBM-licensed program or informal name for a base element of z/OS, that provides for access control by identifying and verifying the users to the system, authorizing access to protected resources, logging the detected unauthorized attempts to enter the system, and logging the detected accesses to protected resources. The formal name for RACF in z/OS is z/OS Security Server.

recovery. The process of rebuilding data after it has been damaged or destroyed, often by using a backup copy of the data or by reapplying transactions recorded in a log.

S

SFF. Single file format

SME. See *system management facilities*

SMS. See *Storage Management Subsystem*

storage administrator. A person in the data processing center who is responsible for defining, implementing, and maintaining storage management policies.

storage class. A collection of storage attributes that identify performance goals and availability requirements, defined by the storage administrator, used to select a device that can meet those goals and requirements.

storage control. The component in a storage subsystem that handles interaction between processor

channel and storage devices, runs channel commands, and controls storage devices.

storage group. A collection of storage volumes and attributes, defined by the storage administrator. The collections can be a group of DASD volumes or tape volumes, or a group of DASD, optical, or tape volumes treated as a single object storage hierarchy. See also *VIO storage group*, *pool storage group*, *tape storage group*, *object storage group*, *object backup storage group*, *dummy storage group*.

storage hierarchy. An arrangement of storage devices with different speeds and capacities. The levels of the storage hierarchy include main storage (memory, DASD cache), primary storage (DASD containing uncompressed data), migration level 1 (DASD containing data in a space-saving format), and migration level 2 (tape cartridges containing data in a space-saving format). See also *primary storage*, *migration level 1*, *migration level 2*, *object storage hierarchy*.

storage management. The activities of data set allocation, placement, monitoring, migration, backup, recall, recovery, and deletion. These can be done either manually or by using automated processes. The Storage Management Subsystem automates these processes for you, while optimizing storage resources. See also *Storage Management Subsystem*.

Storage Management Subsystem (SMS). A DFSMS facility used to automate and centralize the management of storage. Using SMS, a storage administrator describes data allocation characteristics, performance and availability goals, backup and retention requirements, and storage requirements to the system through data class, storage class, management class, storage group, and ACS routine definitions.

storage subsystem. A storage control and its attached storage devices. See also *tape subsystem*.

system-managed storage. Storage managed by the Storage Management Subsystem. SMS attempts to deliver required services for availability, performance, and space to applications. See also *system-managed storage environment*.

DFSMS environment. An environment that helps automate and centralize the management of storage. This is achieved through a combination of hardware, software, and policies. In the DFSMS environment for MVS, this function is provided by DFSMS, DFSORT, and RACF. See also *system-managed storage*.

system-managed tape library. A collection of tape volumes and tape devices, defined in the tape configuration database. A system-managed tape library can be automated or manual. See also *tape library*.

system management facilities (SMF). A component of z/OS that collects input/output (I/O) statistics,

provided at the data set and storage class levels, which helps you monitor the performance of the direct access storage subsystem.

system programmer. A programmer who plans, generates, maintains, extends, and controls the use of an operating system and applications with the aim of improving overall productivity of an installation.

T

tape library. A set of equipment and facilities that support an installation's tape environment. This can include tape storage racks, a set of tape drives, and a set of related tape volumes mounted on those drives. See also *system-managed tape library*.

tape mount management. The methodology used to optimize tape subsystem operation and use, consisting of hardware and software facilities used to manage tape data efficiently.

tape storage group. A type of storage group that contains system-managed private tape volumes. The tape storage group definition specifies the system-managed tape libraries that can contain tape volumes. See also *storage group*.

tape subsystem. A magnetic tape subsystem consisting of a controller and devices, which allows for the storage of user data on tape cartridges. Examples of tape subsystems include the IBM 3490 and 3490E Magnetic Tape Subsystems.

tape volume. A tape volume is the recording space on a single tape cartridge or reel. See also *volume*.

temporary data set. An uncataloged data set whose name begins with & or &&, that is normally used only for the duration of a job or interactive session. Contrast with *permanent data set*.

threshold. A storage group attribute that controls the space usage on DASD volumes, as a percentage of occupied tracks versus total tracks. The *low migration threshold* is used during primary space management and interval migration to determine when to stop processing data. The *high allocation threshold* is used to determine candidate volumes for new data set allocations. Volumes with occupancy lower than the high threshold are selected over volumes that meet or exceed the high threshold value.

time sharing option (TSO). A function of the z/OS operating system that provides interactive time sharing from remote terminals.

TSO. See *time sharing option*

V

virtual input/output (VIO) storage group. A type of storage group that allocates data sets to paging storage, which simulates a DASD volume. VIO storage groups do not contain any actual DASD volumes. See also *storage group*.

virtual storage access method (VSAM). An access method for direct or sequential processing of fixed and variable-length records on direct access devices. The records in a VSAM data set or file can be organized in logical sequence by a key field (key sequence), in the physical sequence in which they are written on the data set or file (entry-sequence), or by relative-record number.

VMA. See *volume mount analyzer*

volume. The storage space on DASD, tape, or optical devices, which is identified by a volume label. See also *DASD volume, optical volume, tape volume*.

volume mount analyzer. A program that helps you analyze your current tape environment. With tape mount management, you can identify data sets that can be redirected to the DASD buffer for management using SMS facilities.

VSAM. See *virtual storage access method*

W

weighted average. DASD space, including free space, which contains not only the data sets generated on primary storage, but also the recalled data sets. During the time of the analysis, this DASD space must be large enough to contain the data sets that will be generated or recalled. The value is calculated over time.

Index

A

abbreviation
 filtering input keyword 157
 GFTAXTR keyword 157
 tailoring output keyword 158

ACCOUNT keyword
 abbreviation 157
 default 157
 primary filter 45
 rules and syntax 64
 using as ACS filter 65
 value in SMF record 45

ACCTFLD keyword
 abbreviation 157
 affecting Top Report 93
 default 157
 rules and syntax 15

ACS (automatic class selection) routine
 ACCOUNT keyword 65
 creating filter list for 41
 DATASET keyword 66
 developing filters for 1, 93
 developing filters with Top Report 93
 EXPDT keyword 68
 JOBNAME keyword 70
 MAXSIZE and MINSIZE keywords 71
 PROGRAM keyword 73
 using GBMAX Report for 89
 using primary filters for 45
 using to redirect data set 2

ACS filter
 creating filter list for 41
 developing with Top Report 93
 identifying 2
 using primary filters for 45

active
 category in Estimate Report 110
 classified by GFTAXTR 23
 selecting with USAGE keyword 78
 tape mount management category 39
 usage pattern in Data Set Report 88
 usage type 36
 using CLASSIFY keyword 136

actual report
 Data Set Report 82
 detailed report 2
 interpreting 82
 Maximum Gigabyte Report 88
 Top Report 92
 Usage Report 95
 Volume Report 101

automated tape library 96

B

backup
 category in Estimate Report 110
 classified by GFTAXTR 23

backup (*continued*)
 selecting with USAGE keyword 78
 tape mount management category 39
 usage pattern in Data Set Report 88
 usage type 36
 using CLASSIFY keyword 136

BCOPY
 classified by GFTAXTR 23
 selecting with USAGE keyword 78
 usage type 36

BLKSIZE keyword
 abbreviation 158
 affecting Data Set Report 82
 affecting Estimate Report 107
 affecting GBMAX Report 89
 affecting IDRC Report 121
 affecting Management Class Report 125
 affecting report 133
 affecting Top Report 93
 affecting Volume Report 102
 default 24, 158
 for migration 24
 rules and syntax 132

BPROGRAM keyword
 abbreviation 158
 affecting Estimate Report 107
 affecting Management Class Report 125
 affecting report 134
 default 158
 rules and syntax 133
 with CLASSIFY 134, 136

buffer
 See tape mount management DASD buffer

C

categorizing data sets
 CLASSIFY keyword 39
 for tape mount management 38
 LARGE keyword 39
 LEVEL0AGE keyword 39
 LEVEL1AGE keyword 39
 order for determining categories 38
 SPLIT keyword 38

CHART keyword
 abbreviation 158
 additional JCL for 50, 135
 default 158
 output example 161
 rules and syntax 134

CLASSIFY keyword
 abbreviation 158
 affecting Estimate Report 107
 affecting Management Class Report 125
 affecting report 136
 categorizing data sets with 39
 default 158

CLASSIFY keyword (*continued*)
 rules and syntax 135
 with BPROGRAM 134, 136
 with LEVEL0AGE 135
 with LEVEL1AGE 135

compaction
 Estimate Report 108
 examining with keyword 69
 flag in SMF record 143
 for tape cartridge 35
 IDRCFACTOR keyword 35

COMPACTION keyword
 abbreviation 158
 affecting Estimate Report 107
 affecting Management Class Report 125
 affecting report 138
 default 158
 rules and syntax 136

continuation character
 for extract program keyword 15
 for filtering input keyword 63
 for tailoring output keyword 131

cost
 Estimate Report summary 107
 of DASD storage 126
 of DASD tracks
 Data Set Report 83
 Management Class Report 129
 of tape mount management DASD buffer 113
 of using tape mount management 107, 117

CUTOFF keyword
 abbreviation 158
 affecting Data Set Report 82
 affecting report 138
 affecting Volume Report 101, 102
 default 158
 rules and syntax 138

D

DASD buffer
 See tape mount management DASD buffer

DASDDEV keyword
 abbreviation 158
 affecting Data Set Report 83
 affecting Estimate Report 107
 affecting Management Class Report 125
 affecting report 139
 affecting Volume Report 102
 controlling device type 88, 105
 default 158
 rules and syntax 138
 using default for Data Set Report 85
 with RPROGRAM 150

data set category
 in Estimate Report 116

data set category (*continued*)
 order for assigning 38

Data Set Report
 actual report 82
 example 86
 GFTASRT3 Report detail 83
 keywords that affect 82
 legend 83
 produced by REPORT keyword 148
 using keywords to generate 82

DATASET keyword
 abbreviation 157
 default 157
 filtering example 46
 primary filter 45
 rules and syntax 65
 using as ACS filter 66

DATE keyword
 abbreviation 157
 affecting Estimate Report 107
 affecting Management Class Report 125
 default 157
 limiting filter 44
 rules and syntax 16, 66
 specifying the end day of sample 18

default
 filtering input keyword 157
 GFTAXTR keyword 157
 tailoring output keyword 158

detailed report
 actual 2, 82
 Data Set Report 82
 Estimate Report 106
 IDRC Report 120
 in output data set 41, 81
 Management Class Report 125
 Maximum Gigabyte Report 88
 optional for a run 1, 81
 simulation 2, 82
 Top Report 92
 Usage Report 95
 using REPORT keyword to generate 148
 volume mount analyzer output 2
 Volume Report 101

DSORT keyword
 abbreviation 158
 affecting Data Set Report 83, 149
 affecting report 140
 default 158
 rules and syntax 139

E

Estimate Report
 data set category 38
 example 108
 keywords that affect 107
 legend 116
 produced by REPORT keyword 148
 savings/cost summary 118
 simulation report 106
 using keywords to generate 106

EXPDT keyword
 abbreviation 157
 default 157

EXPDT keyword (*continued*)
 primary filter 45
 rules and syntax 67
 using as ACS filter 68
 value in SMF record 45

F

FILE keyword
 abbreviation 157
 default 157
 limiting filter 44
 rules and syntax 68
 with MOUNT 72

filter
 automatic class selection (ACS)
 routine 1
 constructing 4
 creating list 41
 developing with Top Report 93
 global data set name character (wild card) 46
 in GFTAVMA input analysis phase 52
 include/exclude sub-keywords 44
 limiting
 DATE 44
 description 44
 FILE 44
 MAXSIZE 44
 MINSIZE 44
 MOUNT 44
 TIME 44
 USAGE 44, 78

primary
 ACCOUNT 45
 DATASET 45
 description 45
 EXPDT 45
 JOBNAME 45
 PROGRAM 45

PROGRAM keyword 93
 running GBMAX Report against 142

secondary
 description 45
 SYSTEMID 45
 UNIT 45
 UNITADDR 45
 VOLUME 45

specifying in GFTAVMA JCL 51
 used for Usage Report 95

filtering
 ACS routine 71
 by usage type 36
 cautions 79
 dropped record 21
 GFTAVMA report 43
 GFTAXTR input phase 18
 keywords for GFTAXTR 18
 keywords for input 63
 MAXSIZE 43
 MAXSIZE and MINSIZE
 keywords 71
 phases of GFTAVMA 42
 SMF data 52
 SPLIT keyword 151
 USAGE 43

filtering (*continued*)
 using Usage Report for 96
 with UNITADDR keyword 77
 with VOLUME keyword 77

FORCECAP keyword
 abbreviation 158
 affecting Data Set Report 83, 85
 affecting Estimate Report 107
 affecting GBMAX Report 89
 affecting GFTASRT2 Report 55
 affecting IDRC Report 121
 affecting Management Class Report 125
 affecting report 141
 affecting Volume Report 102
 default 158
 rules and syntax 140

FSPACE keyword
 abbreviation 158
 affecting Estimate Report 107
 affecting report 142
 default 158
 rules and syntax 141
 using to change free space 116

G

GBMAX (Maximum Gigabyte) Report
 actual report 88
 example 90
 keywords that affect 89
 produced by REPORT keyword 148
 running against ACS filter list 142
 using keywords to generate 88

GDG keyword
 abbreviation 158
 affecting Data Set Report 83
 affecting Estimate Report 107
 affecting IDRC Report 120
 affecting report 142
 default 158
 rules and syntax 142

GFTASRT1 Report
 example 52
 order of reports 42

GFTASRT2 Report
 order of reports 42
 sample 53, 55
 supplemented by Volume Report 102

GFTASRT3 Report
 order of reports 42
 sample 57, 59
 supplemented by Data Set Report 83

GFTAVMAP JCL member 43, 48
 GFTAXTRP JCL member 11
 global data set name character (wild card) 46, 76

H

HSM
 category in Estimate Report 110
 selecting with USAGE keyword 78
 tape mount management category 88
 usage type 36

I

- IDRC Report
 - data set category 38
 - detailed report 120
 - example 123
 - IDRCFACTOR keyword 144
 - keywords that affect 121
 - legend 122
 - produced by REPORT keyword 148
 - simulation report 120
 - TAPEDEV keyword 144
 - TAPELEN keyword 144
 - using keywords to generate 120
- IDRCFACTOR keyword
 - abbreviation 158
 - affecting Estimate Report 107
 - affecting IDRC Report 121
 - affecting Management Class Report 126
 - affecting report 144
 - calculating IDRC Report 122
 - compaction for tape cartridge 35
 - controlling compaction 144
 - default 158
 - rules and syntax 143
 - using if no specified compaction 122
 - with FORCECAP 141
- IFASMFDP dump utility 12, 13
- include/exclude filter
 - category 44
 - definition 44
 - example 46
 - syntax 45
- input
 - cannot merge or subset 42
 - creating data set with GFTAXTR 9
 - filtering 44
 - GFTASRT1 processing phase 52
 - GFTAVMAP JCL member 51
 - keywords that filter 63
 - processing data with GFTAXTR 9
 - required SMF record 10
 - to GFTAVMA 18, 41
 - to GFTAXTR 10
- interval migration
 - clearing DASD buffer with 120
 - hourly 117
 - to SMS storage group 2

J

- JCL
 - additional
 - needed for CHART keyword 50, 135
 - needed for SPLIT keyword 50, 152
 - GFTAVMAP member
 - example 49
 - submitting 43, 48
 - using keywords in 43
 - GFTAXTRP member
 - example 11
 - submitting 11
 - using keywords in 14
 - specifying keywords in 14, 51

- JOBNAME keyword
 - abbreviation 157
 - default 157
 - filtering example 47
 - primary filter 45
 - rules and syntax 69
 - using as ACS filter 70
 - value in SMF record 45

K

- keyword
 - affecting Data Set Report 82
 - affecting Estimate Report 107
 - affecting GMAX Report 89
 - affecting IDRC Report 121
 - affecting Management Class Report 125
 - affecting Top Report 93
 - affecting Volume Report 102
- extract program (GFTAXTR)
 - abbreviation 157
 - ACCTFLD 15
 - continuation character 15
 - DATE 16
 - default 157
 - LINES 17
 - TIME 17
 - using 15
- filtering input 63
 - abbreviation 157
 - ACCOUNT 64
 - comments in 63
 - continuation character 63
 - DATASET 65
 - DATE 66
 - default 157
 - EXPDT 67
 - FILE 68, 72
 - JOBNAME 69
 - MAXSIZE 70
 - MINSIZE 70
 - MOUNT 69, 72
 - PROGRAM 73
 - rule 63
 - SYSTEMID 74
 - TIME 74
 - UNIT 75
 - UNITADDR 76
 - USAGE 36, 78
 - using 63
 - VOLUME 79
- generating Usage Report 95
- specifying in GFTAVMA JCL 51
- specifying in GFTAXTR JCL 14
- tailoring output 131
 - abbreviation 158
 - BLKSIZE 132
 - BPROGRAM 133, 136
 - CHART 134
 - CLASSIFY 134, 135
 - comments in 131
 - COMPACTION 136
 - continuation character 131
 - CUTOFF 138
 - DASDDEV 138, 150
 - default 158

- keyword (*continued*)
 - tailoring output (*continued*)
 - DSORT 139
 - FORCECAP 140
 - FSPACE 141
 - GDG 142
 - IDRCFACTOR 143
 - L1FSPACE 147
 - LARGE 144
 - LEVEL0AGE 145
 - LEVEL1AGE 146
 - LINES 147
 - NOFORCECAP 140
 - REPORT 142, 148
 - RPROGRAM 150
 - rule 131
 - SPLIT 151
 - TAPEDEV 143, 152
 - TAPELEN 143, 153
 - using 131
 - VSORT 154

L

- L1FSPACE keyword
 - abbreviation 158
 - affecting Estimate Report 107
 - affecting report 148
 - default 158
 - rules and syntax 147
 - using to change free space 116
- large
 - category in Estimate Report 110
 - tape mount management category 39
- LARGE keyword
 - abbreviation 158
 - affecting Estimate Report 107
 - affecting IDRC Report 121
 - affecting Management Class Report 126
 - affecting report 145
 - categorizing data sets with 39
 - default 158
 - in target environment 110
 - rules and syntax 144
 - selecting data sets with 121
 - setting value for IDRC Report 120
 - using for Top Report 93
- LEVEL0AGE keyword
 - abbreviation 158
 - affecting Estimate Report 107
 - affecting Management Class Report 125
 - affecting report 146
 - categorizing data sets with 39
 - default 158
 - rules and syntax 145
 - with CLASSIFY 135
- LEVEL1AGE keyword
 - abbreviation 158
 - affecting Estimate Report 107
 - affecting Management Class Report 125
 - affecting report 147
 - categorizing data sets with 39
 - default 158
 - rules and syntax 146

LEVEL1AGE keyword (*continued*)
with CLASSIFY 135
LINES keyword
abbreviation 157, 158
default 157, 158
extract program (GFTAXTR) 17
rules and syntax 147

M

Management Class Report
data set category 38
example 127
keywords that affect 125
produced by REPORT keyword 148
simulation report 125
using keywords to generate 125
MAXSIZE keyword
abbreviation 157
affecting Data Set Report 82
default 157
limiting filter 44
rules and syntax 70
MINSIZE keyword
abbreviation 157
affecting Data Set Report 82
default 157
limiting filter 44
rules and syntax 70
MOUNT keyword
abbreviation 157
default 157
limiting filter 44
rules and syntax 72
with FILE 69

N

NOFORCECAP keyword
abbreviation 158
affecting report 141
default 158
rules and syntax 140

O

order to request report 43, 51
over-filtering 28, 51

P

primary filter 45
problems with output
for GFTAVMA 51
for GFTAXTR 28
over-filtering 51
PROGRAM keyword
abbreviation 157
affecting Top Report 93
default 157
primary filter 45
rules and syntax 73
using as ACS filter 73, 93
value in SMF record 45

R

report
automatic for every run
description 81
GFTASRT1 52
GFTASRT2 53
GFTASRT3 57
optional for a run
Data Set 82
description 81
Estimate 106
GBMAX (Maximum Gigabyte) 88
IDRC 120
Management Class 125
Top 92
Usage 95
Volume 101
REPORT keyword
abbreviation 158
default 158
generating Estimate Report 117
rules and syntax 148
sub-keyword
DATASET 148
ESTIMATE 148
GBMAX 148
IDRC 148
MGMTCLAS 148
TOP 148
USAGE 149
VOLUME 149
RPROGRAM keyword
abbreviation 158
affecting Data Set Report 83
affecting Estimate Report 107
affecting GBMAX Report 89
affecting IDRC Report 121
affecting Management Class
Report 126
affecting report 151
affecting Volume Report 102
default 158
rules and syntax 150

S

secondary filter 45
simulation report
detailed report 2
Estimate Report 106
IDRC Report 120
interpreting 106
Management Class Report 125
single
selecting with USAGE keyword 78
usage type 36
SMF record
analyzing 1
collecting 4
description of type 11
end sample on busy tape day 14
flag for IDRC compaction 143
input to GFTAXTR 10
merging for multiple systems 6
one month period 5
over-filtering 51

SMF record (*continued*)
processing with GFTAXTR 9
re-created information 10
representing tape cartridge 35
required type 10
selected with keyword 63
type 14 133
type 15 133
using for analysis 4
valid record level 4
SMFIN
in GFTAXTR JCL 13
input to GFTAXTR 11
SMS (Storage Management Subsystem)
directing data sets 3
storage group 2
SPLIT keyword
abbreviation 158
additional JCL for 50, 152
categorizing data sets with 38
default 158
filtering 151
rules and syntax 151
storage group
tape mount management DASD
buffer 2, 121
summary report
automatic for every run 1, 81
date and time 26
GFTASRT1 52
GFTASRT2 53, 55
GFTASRT3 57, 59
if REPORT not specified 149
in output data set 41, 81
interpreting 51
volume mount analyzer output 2
syntax
filtering input keyword 63, 159
GFTAXTR keyword 15, 159
include/exclude filter 45
tailoring output keyword 131, 159
SYSTEMID keyword
abbreviation 157
default 157
rules and syntax 74
secondary filter 45

T

tape mount management
benefits of 3, 121
categorizing data sets 38
cost and savings 107, 117
data set category 38
defining category 116
determining candidate for 41
filling tape 121
implementing 109, 112
reducing mounts with 61
type of tape cartridge 35
using DFSMSshsm 2
using Estimate Report for 116, 148
using GBMAX Report for 89
using LARGE keyword for 121
using primary filters for 45
using the volume mount analyzer
for 2

tape mount management (*continued*)
 using Top Report for 93

tape mount management DASD buffer
 calculating size
 for tape mount management 1
 using FSPACE 141
 using GBMAX Report 89
 using LIFSPACE 147

clearing with DFSMSHsm 120

cost breakdown in Estimate
 Report 113

determining cost of 117, 120

level 2 not part of 115

storage group 2, 121

temporary storage 3

using Estimate Report 117, 120

using GBMAX Report 90

using Top Report 93

TAPEDEV keyword
 abbreviation 158

affecting Estimate Report 107, 118

affecting IDRC Report 121

affecting Management Class
 Report 126

affecting report 153

calculating IDRC Report 122

controlling device type 144

default 158

rules and syntax 152

specifying hardware 110

with IDRCFACTOR keyword 143

with TAPELEN keyword 153

TAPELEN keyword
 abbreviation 158

affecting Estimate Report 107, 118

affecting IDRC Report 121

affecting Management Class
 Report 126

affecting report 154

calculating IDRC Report 122

controlling device type 144

default 158

rules and syntax 153

specifying hardware 110

with IDRCFACTOR keyword 143

with TAPEDEV keyword 154

temporary
 category in Estimate Report 110

selecting with USAGE keyword 78

tape mount management category 39

usage type 36

TIME keyword
 abbreviation 157

affecting Estimate Report 107

affecting Management Class
 Report 125

default 157

extract program (GFTAXTR) 17

limiting filter 44

rules and syntax 74

Top Report
 actual report 92

default 149

example 94

keywords that affect 93

parameter 149

produced by REPORT keyword 148

Top Report (*continued*)
 types generated by parameter
 accounting information 149

days since last reference 149

expiration date 149

high level qualifier 149

job name 149

last reference age 149

low level qualifier 149

program name 149

using keywords to generate 92

using PROGRAM as ACS filter 93

using to determine filters 93

TSO SUBMIT command 11, 48

U

UNIT keyword
 abbreviation 157

default 157

rules and syntax 75

secondary filter 45

with UNITADDR keyword 76

UNITADDR keyword
 abbreviation 157

default 157

rules and syntax 76

secondary filter 45

with UNIT keyword 76

with VOLUME keyword 77, 79

USAGE keyword
 abbreviation 157

default 157

filtering input 36

limiting filter 44

rules and syntax 78

Usage Report
 actual report 95

example 97

produced by REPORT keyword 149

shared system considerations 5

using keywords to generate 95

when generated 42

usage type
 description 36

order for assigning 36

V

VOLUME keyword
 abbreviation 158

default 158

rules and syntax 79

secondary filter 45

with UNITADDR keyword 77, 79

Volume Report
 actual report 101

example 104

GFTASRT2 Report detail 102

keywords that affect 102

legend 103

produced by REPORT keyword 149

using keywords to generate 101

when generated 42

VSORT keyword
 abbreviation 158

VSORT keyword (*continued*)
 affecting report 155

affecting Volume Report 102, 149

default 158

rules and syntax 154

W

wild card
 See global data set name character

X

XTRCIN
 DD statement 13, 42

in GFTAXTRP JCL member 18

output of GFTAXTR 11, 51

total SMF records in 21

XTRIN
 DD statement 41

in GFTAVMAP JCL 49

input to GFTAVMA 41

output of GFTAXTR 41

update for SPLIT 50

Readers' Comments — We'd Like to Hear from You

z/OS

DFSMS: Using the Volume Mount Analyzer

Publication No. SC26-7413-01

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department 55JA, Mail Station P384
2455 South Road
Poughkeepsie, NY 12601-5400



Fold and Tape

Please do not staple

Fold and Tape



Program Number: 5694-A01

Printed in USA

SC26-7413-01

