

CA MICS® Resource Management

StorageMate Option Guide

Release 12.9



This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by CA at any time. This Documentation is proprietary information of CA and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA.

If you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2014 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At <http://ca.com/support>, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at <http://ca.com/docs>.

Contents

Chapter 1: OVERVIEW	9
1.1 Primary Areas of Application.....	11
1.2 Reporting and Inquiry Facilities.....	18
1.3 Requirements.....	19
 Chapter 2: ACCESSING CA MICS StoageMate	 23
2.1 Security Authorization.....	23
2.2 Database Information	24
2.3 Levels of Summarization	26
2.4 CA MICS Update Schedule.....	27
 Chapter 3: CUSTOMIZING CA MICS StorageMate	 29
3.1 Defining Data Sets and Options.....	30
3.2 Establish Global Processing Options	34
3.3 Executing an Inquiry	39
3.4 Volume Group Tables.....	46
3.4.1 Creating a Volume Group Table	50
3.4.2 Modifying a Volume Group Table	52
3.5 Data Set Group Tables.....	53
3.5.1 Creating a Data Set Ownership Table	55
3.5.2 Modifying a Data Set Ownership Table.....	57
3.6 Bypassing the Execution Run Time Screen.....	59
3.7 Setting Up Your JCL Defaults	60
3.8 Setting Up Your Graphic Defaults.....	61
3.9 Data Filtering	63
3.10 How to Enable Neugents technology in CA MICS.....	67
3.10.1 Requirements.....	67
3.10.2 Activate the CA MICS Root Directory	67
3.10.3 Enable Neugents technology	68
 Chapter 4: REPORTS	 69
4.1 SMS Applications	71
4.2 Non-SMS Applications.....	76
4.3 Report Descriptions.....	79
ACS Routine Definition / By Pattern (STGELB)	80

ACS Routine Definition / By Group (STGELB)	84
ACS Routine Definition / Detail (STGELB)	87
ACS Test Case Creation (STGELF)	91
ACS Testing Results / Statistics (STGELG)	94
ACS Testing Results / Detail (STGELG)	98
ACS Testing Results / Exceptions (STGELG)	102
ACS Testing Results / Analysis (STGELG)	106
ACS Testing Results / Group (STGELG)	110
ACS Testing Results / Owner (STGELG)	114
Allocation Problems / Chart (STGECB)	118
Allocation Problems / Summary (STGECB)	121
Allocation Problems / Detail (STGECB)	124
ASTEX / Top Resource Consumers / Chart (STGEPF)	128
ASTEX / Top Resource Consumers / Summary (STGEPF)	135
ASTEX / Top Resource Consumers / Detail (STGEPF)	140
Daily Application Resource Usage / Plot (STGERG)	148
Daily Application Resource Usage / Summary (STGERG)	157
Daily ASTEX Cache Statistics / Plot (STGEPE)	165
Daily ASTEX Cache Statistics / Summary (STGEPE)	172
Daily DASD Resource Usage / Plot (STGEJE)	178
Daily DASD Resource Usage / Summary (STGEJE)	186
Daily HSM DASD Usage / Chart (STGEJC)	193
Daily HSM DASD Usage / Summary (STGEJC)	199
DASD-to-Tape Candidates (STGERA)	205
DASD/HSM Data Movement / Plot (STGECE)	208
DASD/HSM Data Movement / Summary (STGECE)	211
DASD/HSM Data Movement / History (STGECE)	218
DASD/HSM Usage Snapshot / Chart (STGERF)	222
DASD/HSM Usage Snapshot / Summary (STGERF)	226
Data Set Names (STGEFA)	230
Data Set Performance / Chart (STGEPA)	234
Data Set Performance / Summary (STGEPA)	237
Deletion/Migrati Thresholds / Chart (STGEJA)	240
Deletion/Migrati Thresholds / Summary (STGEJA)	244
GDG Management and Retention / By Prefix (STGEFE)	250
GDG Management and Retention / By Base Name (STGEFE)	254
GDG Management and Retention / Detail (STGEFE)	259
I/O Activity by Workload / Chart (STGEPC)	262
I/O Activity by Workload / Detail (STGEPC)	265
Monthly Application Resource Usage / Plot (STGERE)	269
Monthly Application Resource Usage / Summary (SEGERE)	278
Monthly DASD Resource Usage / Plot (STGERH)	286

Monthly DASD Resource Usage / Summary (STGERH).....	294
New Data Set Allocation / Chart (STGECG)	301
New Data Set Allocation / Summary (STGECG).....	305
New Data Set Allocation / Detail (STGECG).....	308
Non-VSAM Attributes / Summary (STGEFB)	313
Non-VSAM Attributes / Chart (STGEFB)	317
Non-VSAM Attributes / Node Analysis (STGEFB)	320
Non-VSAM Attributes / Detail (STGEFB)	323
Non-VSAM Wasted Space / Chart (STGECC)	327
Non-VSAM Wasted Space / Summary (STGECC)	329
Non-VSAM Wasted Space / Detail (STGECC).....	333
SMS Class Usage / Chart (STGELE)	337
SMS Class Usage / Summary (STGELE)	340
SMS Class Usage / Owner (STGELE)	344
SMS Class Usage / Detail (STGELE).....	349
SMS Class Usage / Descriptive Statistics (STGELE)	353
SMS Class Usage / Distribution of Allocated Space (STGELE).....	356
SMS Class Usage / Analysis of Space Usage (STGELE)	359
SMS Implementation / By Owner (STGELC)	363
SMS Implementation / by Volume (STGELC).....	366
Space Allocation by User / Chart (STGERC).....	369
Space Allocation by User / Detail (STGERC)	372
Space Management Function Simulation / Chart (STGEJB)	375
Space Management Function Simulation / Summary (STGEJB).....	378
Space Management Function Simulation / Detail (STGEJB).....	382
Tape-to-DASD Candidates (STGERB)	385
Unmanageable Data Sets / Summary (STGELA)	388
Unmanageable Data Sets / Detail (STGELA).....	392
Volume Configuration / Summary (STGEPB).....	397
Volume Configuration / Detail (STGEPB)	400
Analysis of I/O Rate and Response Time (STGEPB)	404
Listing of All Exception Volumes by Cluster (STGEPB)	407
Volume Group Activity by Workload / Chart (STGEHB)	410
Volume Group Allocation by User / Chart (STGEHC)	413
Volume Group Allocation by User / Summary (STGEHC)	416
Volume Group Configuration / Totals (STGEHA).....	420
Volume Group Configuration / Detail (STGEHA)	423
Volume Problems (STGECA).....	427
VSAM Attributes / Summary (STGEFC)	433
VSAM Attributes / Chart (STGEFC)	437
VSAM Attributes / Node Analysis (STGEFC)	440
VSAM Attributes / Detail (STGEFC)	443

VSAM Wasted Space / Chart (STGECECF)	446
VSAM Wasted Space / Summary (STGECECF)	449
VSAM Wasted Space / Detail (STGECECF)	453
VSAM Wasted Space / Analysis (STGECECF)	456
FACILITY PROCEDURES	463
Appendix A: ERROR MESSAGES	469
Messages	469
Appendix B: INSTALLATION	517
Appendix C: TAPE INTERFACE	521
Appendix D: CASE STUDIES	523
Appendix E: NON-VSAM DATA SET ACTION-RELATED INTERFACE	541
Appendix F: VSAM DATASET ACTIVITY-RELATED INTERFACE	545

Chapter 1: OVERVIEW

The amount of information retained in data centers on magnetic or electronic media is growing at an astronomical rate. There are two major reasons for this growth. As the number of online applications increases, so does the amount of data available to the interactive user. Also, new technology in the image processing area is adding to the amount of data that needs to be stored online.

As a result, your role as storage administrator has become more critical to the success of your installation. Complicating matters, the online availability of mission critical applications limits the amount of downtime available to you in which to perform the necessary maintenance on storage resources.

The StorageMate product allows you to identify problems in a timely manner, identify and verify solutions, assist in storage group planning, and simulate the effects of proposed changes on the storage configuration. The product offers a number of facilities that enable you to perform these functions, not just generate reports.

As shown in Figure 1-1, the CA MICS StorageMate product uses the data integration and processing capabilities of the CA MICS DASD Space Collector, CA MICS DASD Space Analyzer, CA MICS MVS Batch and Operations Analyzer, CA MICS MVS Hardware and SCP Analyzer, and CA MICS Application Extension for ASTEX. Data from IBM's Data Facility Hierarchical Storage Manager (DFHSM) can also be processed if you have implemented the HSM Inventory option of the CA MICS DASD Space Analyzer. These products populate information areas in the CA MICS database that StorageMate then uses. Information about your tape library can also be read from an external source for some reports. The StorageMate product is an online interactive tool you can use to access and analyze information and produce results via graphic or tabular reports.




Figure 1-1. CA MICS StorageMate Data Flow

This section contains the following topics:

- [1.1 Primary Areas of Application](#) (see page 11)
- [1.2 Reporting and Inquiry Facilities](#) (see page 18)
- [1.3 Requirements](#) (see page 19)

1.1 Primary Areas of Application

The StorageMate product can be used in a number of areas the storage administrator typically will encounter. The primary areas of application are discussed below.

Problem Management

The CA MICS StorageMate product can help identify existing and potential problems at both a detailed and summary level that reduce the effectiveness of your storage resources.

At the volume level, volumes without indexed VTOCs that could cause degraded response time are easily identified. Other potential volume errors that CA MICS StorageMate can identify include VTOC structure damage, a missing or undersized VSAM Volume Data Set (VVDS), and excessive seek delays caused by improper data set placement.

CA MICS StorageMate identifies allocation errors by scanning the SMF job files in the CA MICS database and looking for space-related abends. The failing jobs and the amount of time and resources expended in recovering from the failure are identified. Assessment of your effectiveness is often tied to the number of these abends occurring within your installation. This report can establish or monitor a service level agreement between you and the data center management team.

Another facility in the StorageMate product identifies data sets that are allocated so that they cannot easily be managed by automated storage management software. For example, products such as DFHSM and DFSMS require that the data sets they manage be cataloged. Uncataloged data sets must be manipulated manually. The ability to identify these problem data sets will reduce your duties and help facilitate the implementation of automated products. The Unmanageable Data Sets facility within CA MICS StorageMate that detects these errors can also generate a jobstream that will correct some of the errors encountered.

Large portions of storage resources are often wasted by poor allocation practices. Most of this wasted space is caused by one of the following conditions:

- o The user has allocated far more space than was actually needed.
- o The user is storing data using a block size that is not

efficient for the device on which the data is stored.

- o The user has allocated VSAM data sets using options that are not efficient for the storage device or for the application that accesses the data.

All of these conditions can be detected using the Non-VSAM Wasted Space and the VSAM Wasted Space facilities included in the StorageMate product. Reports within these facilities detect these problem data sets and, optionally, show you only the worst violators. This allows you to concentrate on correcting those data sets that are wasting the most space. Another option of these facilities shows wasted space at the department or group level, to help you quickly identify groups that are not using storage effectively.

Another facility provided by CA MICS StorageMate, the DASD/HSM Data Movement facility, allows you to monitor the movement of data between your primary DASD volumes and DFHSM. Many installations use DFHSM to move unused data sets to DFHSM migration storage, thus freeing their critical, primary DASD space for the allocation of new data sets. This process needs to be monitored to ensure that DFHSM parameters and SMS management class values are still working as expected. Even a well-tuned DFHSM system may not function well if your data set access patterns change over time. Intentional abuse of primary space or DFHSM space by users or groups can also be detected with these reports.

One of the critical activities you should be performing regularly is the monitoring of new data sets. Early detection of excessive space consumption will help prevent short-term performance problems and long-term capacity problems. Even users of DFSMS need to perform this important monitoring, as users under DFSMS may still have the ability to override system defaults and change the size and location of their new data sets. The New Data Set Allocation facility within StorageMate allows you to regularly monitor the allocation patterns of new data sets. New allocation activity can be summarized according to volume, DFSMS construct, DFSMS Storage Group or application area. Patterns present on the summary reports can be further investigated using detailed reporting.

The following reports can assist you with problem management activities:

- o Unmanageable Data Sets
- o Volume Problems

- o Allocation Problems
- o New Data Set Allocation
- o Non-VSAM Wasted Space
- o VSAM Wasted Space
- o DASD/HSM Data Movement

For more information on these reports, see Section 4.3 of this guide.

Data Set Standards

The StorageMate product helps you establish and enforce data set standards, that, in turn helps you simplify operations and better control your storage environment.

Data set name standardization allows for usage and service requirements to be determined by a simple review of the data set name. The Data Set Names facility of CA MICS StorageMate scans the CA MICS database and reports the frequency of usage of data set name nodes. It also flags data set name exceptions, thereby facilitating enforcement of existing data set standards.

Two facilities, the VSAM Attributes facility and the Non-VSAM Attributes facility within the StorageMate product, allow you to quickly identify common allocation patterns within your data sets, which allows you to detect groups of data sets that are not using storage resources efficiently. Once these groups are identified, additional options allow you to investigate the problem groups in greater detail.

The use of Generation Data Groups (GDGs) is helpful to users, but can be a burden to you. The GDG Management and Retention facility within CA MICS StorageMate allows you to summarize the use of GDGs. The reports produced from this facility allow you to identify potential problems within these groups, such as groups allocated on too few volumes for adequate recovery or duplicate or missing generations.

A simulation option allows you to determine the effects of proposed GDG options without actually making any changes or moving any data.

The following reports can assist you with data set standards:

- o Data Set Names
- o Non-VSAM Attributes
- o VSAM Attributes

- o GDG Management and Retention

For more information on these reports, refer to Section 4.3 of this guide.

Volume Grouping

One recent trend in storage management is the concept of pooling similar DASD volumes into a logical entity. These logical groups are often referred to as volume pools, volume groups, or (for SMS users) storage groups. This trend is encouraged by the rapid increase in the amount of data that must be managed and also by the appearance of automated storage management products that expect or encourage volume grouping.

While volume grouping allows you more simplicity and better control of the environment, it also complicates the tasks of reporting and monitoring because these are elevated from the volume level to the volume group level.

The Volume Group Configuration, Volume Group Activity by Workload, and Volume Group Allocations by User facilities within CA MICS StorageMate simplify reporting at the volume group level. They analyze the allocation, performance, and activity of each volume group. This allows you to track service level agreements at the volume group level and detect the possible misuse of groups in certain environments.

A simulation option within each of these facilities allows easy testing of proposed changes in your volume group definitions. This allows you to experiment with alternate configurations and determine the optimum groupings without causing any adverse effect on your users.

The following reports can help you with volume grouping:

- o Volume Group Configuration
- o Volume Group Activity by Workload
- o Volume Group Allocations by User

For more information on these reports, refer to Section 4.3 of this guide.

Space Availability

Space availability has historically been a concern of storage administrators because users always expect to have plenty of space at their disposal. The need for this level of service

has led to the development of automated products such as DFHSM, DFSMS, and DMS/OS that have made the Storage Administrator's job easier.

The StorageMate product provides several tools that help you establish and tune many of the thresholds you must define to products such as DFHSM. One analysis helps you to determine the frequency patterns of data set usage and to estimate the values that should be used to drive your automated products at optimum efficiency. CA MICS StorageMate also helps you verify that your automated space management tools (DFSMS and DMS/OS) are working properly and set up or modify the migration, deletion, and release thresholds used by these products. When threshold changes are required, the Space Management Function Simulation facility allows you to determine the effects of those changes before any data is actually affected.

Even a well tuned storage system can experience availability problems if it is not monitored and adjusted regularly. Application changes and normal business cycles are only two of many events that may change storage availability patterns. The Daily DASD Resource Usage and Daily Application Resource Usage facilities can be used to monitor space availability and take corrective action before problems develop. For DFHSM users, the Daily HSM DASD Usage facility provides the storage picture from the viewpoint of HSM. This not only provides early warning for potential space problems, but also allows you to measure the effectiveness of HSM options and space management activities.

The following reports can help you with space availability:

- o Deletion/Migration/Release Thresholds
- o Space Management Function Simulation
- o Daily DASD Resource Usage
- o Daily Application Resource Usage
- o Daily HSM DASD Usage

For more information on these reports, refer to Section 4.3 of this guide.

SMS Migration and Monitoring

DFSMS, DMS/OS, and other automated products have revolutionized storage administration. These automated tools allow each storage administrator to manage more megabytes, data sets, and physical volumes.

These tools must be properly implemented and tuned if they are to do the job correctly. The key is proper planning and preparation.

The CA MICS StorageMate product can give you the proper level of knowledge about the data center and provide data useful in deciding how the automated tools should operate. CA MICS StorageMate can also monitor your SMS implementation progress so that management stays informed of the status of the migration effort.

After DFSMS or other automated products are fully implemented, additional StorageMate facilities can be used to monitor the efforts of these automated products. This helps you ensure the products are operating as designed and quickly identifies user or system errors that could affect the performance and integrity of the automated system.

The following reports can help you with SMS migration and monitoring:

- o ACS Routine Definition
- o ACS Test Case Creation
- o ACS Testing Results
- o SMS Implementation
- o SMS Class Usage

For more information on these reports, refer to Section 4.3 of this guide.

Performance Management

Information that is stored electronically or magnetically must be accessible in a timely manner. One of the strategic challenges you face is maintaining the performance of data sets, volumes, volume pools, and other storage entities. CA MICS StorageMate detects performance problems before they become critical and affect your storage users.

While some performance problems can be detected from RMF data, more specific information about many performance problems must be obtained from a more detailed storage monitor such as ASTEX. Some CA MICS StorageMate reports read information supplied by the CA MICS Application Extension for ASTEX. These reports often give very precise information about the causes of and solutions to performance problems.

The following reports can help you with performance management:

- o Data Set Performance
- o Volume Configuration
- o I/O Activity by Workload
- o Daily ASTEX Cache Statistics
- o ASTEX / Top Resource Consumers

For more information on these reports, refer to Section 4.3 of this guide.

Capacity Management

The growth of critical online systems and expansion of data across distributed networks has caused a growth in storage requirements similar to the growth that has occurred in processing power over the last two decades. You must constantly identify and track resource requirements for critical applications and project the future growth of storage resource requirements so that hardware upgrades or data management plans can be implemented before any shortage occurs.

The StorageMate product allows you to identify and monitor consumption of critical storage resources at many different levels.

The following reports can help you with capacity management:

- o DASD-to-Tape Candidates
- o Tape-to-DASD Candidates
- o Space Allocation by User
- o Monthly DASD Resource Usage
- o Monthly Application Resource Usage
- o DASD/HSM Usage Snapshot

For more information on these reports, refer to Section 4.3 of this guide.

1.2 Reporting and Inquiry Facilities

The StorageMate product provides a number of facilities that can produce summary, detail, and other reports. These reports encompass several aspects of storage management, including problem management, data set standards, volume grouping, space availability, SMS migration and monitoring, performance management, and capacity management. The names of the reports that fall into these functional categories are provided in Section 1.1.

CA MICS StorageMate uses the CA MICS Information Center Facility (MICF), an interactive inquiry facility of CA MICS, to produce the reports in either interactive or batch mode. The reports present the information in tabular and graphic formats that can be displayed at a terminal, cataloged for later viewing, or printed on a hard copy device.

Execution of StorageMate reports under the MICF Production Reporting facility is supported. Because reports executing in that facility have no method for obtaining options at execution time, all StorageMate reports are distributed with option defaults that will be used when running under the MICF Production Reporting facility. Generally, these defaults tend to request summary reports rather than detail reports, and provide no special data filtering in terms of selected sysids, volumes, etc. For some reports, installation defaults must be provided. These defaults are set in the \$STGEXIT options member, and are discussed in Section 3.2 of this guide.

More detailed information about the reports and facilities is available in Chapter 4 of this guide.

1.3 Requirements

The CA MICS StorageMate product requires that the following CA products are installed:

- o CA MICS Platform
- o CA MICS MVS Hardware and SCP Analyzer
- o CA MICS MVS Batch and Operations Analyzer
- o CA MICS DASD Space Collector
- o CA MICS DASD Space Analyzer

The following products and options will extend the reporting capabilities of the StorageMate product. If these products are not installed or implemented, some reports and report options may not be available:

- o CA MICS Application Extension for ASTEX
- o HSM Inventory option for the CA MICS DASD Space Analyzer

Data Sources

As illustrated in Figure 1-1, the CA MICS database contains many information areas. CA MICS StorageMate uses the VCA, SCP, HAR, BAT, HSM, HSB, and AST information areas of the CA MICS database. Information from these areas is integrated to give you a complete view of the storage environment.

VTOC/CATALOG ACTIVITY (VCA) INFORMATION AREA

The VCA information area has files that are populated from the CA MICS DASD Space Analyzer (VCA) and DASD Space Collector (VCC). VCC is usually run once a day on each processing complex. As the name suggests, it scans DASD VTOCs and MVS system catalogs, extracting information about data sets and volumes. This information is analyzed and summarized into a number of files.

HARDWARE (HAR) AND SYSTEM CONTROL PROGRAM (SCP) INFORMATION AREAS

The HAR information area has files that are populated from MVS RMF (Resource Measurement Facility) records via the CA MICS Hardware and SCP Analyzer. These raw records are written by RMF at a user-specified interval. Most installations use intervals such as 15, 20, or 30 minutes as the standard. The information from this area reported by CA MICS StorageMate includes device response times and I/O rates. CA MICS StorageMate also uses the SCP information

area in reporting.

BATCH (BAT) INFORMATION AREA

The Batch Information Area contains information populated from MVS SMF (System Management Facility) records via the CA MICS Batch and Operations Analyzer. These records are written when certain events occur, such as the end of a job step. Information from these files used by CA MICS StorageMate includes device I/O activity by MVS workload and information about work that was delayed due to storage-related failures.

HSM MIGRATION (HSM) AND BACKUP (HSB) INFORMATION AREAS

These information areas contain information about user data sets that are under control of DFHSM. To use the HSM information area, you must have implemented the HSM Inventory option of the CA MICS DASD Space Analyzer. This option causes the DASD Space Collector to interface with DFHSM during daily space collection, and collect and write records containing information about data sets migrated or backed up by DFHSM. These records are then processed on a daily basis by the analyzer and are used to create the files in the HSM information areas.

ASTEX (AST) INFORMATION AREA

The ASTEX information area contains DASD and cache performance data supplied by the CA Automated STorage EXpert (ASTEX) product. To use this data, you must have installed the CA MICS Application Extension for ASTEX. The data provided in this information area contains very detailed performance information collected at the data set, volume, storage group, and control unit level. ASTEX users should review their ASTEX options carefully to make sure the intervals and recording levels chosen for ASTEX reporting are granular enough to provide adequate detail in this information area.

TAPE DATA (EXTERNAL)

You can use an external interface to process data that resides outside the CA MICS database. This interface allows you to extract information related to your automated tape library, format that data set into a SAS database, and include that database in CA MICS StorageMate reports. While this interface should work with most common tape management products, samples are provided that are designed to support

the most commonly used tape library management products. See Appendix C for a further discussion of this facility.

TIMESPANS AND FILES

The chart below displays the files and timespans used by CA MICS StorageMate in each information area. Depending on the options or defaults used when CA MICS was installed, some of these files may not be active in the required timespans. In this case, you will not be able to use some reports until those files are activated. Check with your CA MICS Administrator to make sure all needed files are active in the required timespans.

Information		
Area	File	Timespans

AST	ADS	DETAIL
AST	AVO	DAYS
BAT	_JS	DETAIL
BAT	JOB	DETAIL
BAT	PGM	DETAIL
BAT	WDA	DAYS
HAR	DVA	DAYS
HSB	BAC	DETAIL, DAYS, MONTHS
HSM	MIG	DETAIL, DAYS, MONTHS
HSM	VOL	DETAIL
SCP	PGA	DAYS
VCA	_VS	DETAIL, DAYS, MONTHS
VCA	DAA	DETAIL, DAYS, MONTHS
VCA	VOA	DETAIL, MONTHS

Chapter 2: ACCESSING CA MICS StorageMate

CA MICS StorageMate users may need assistance from their CA MICS System Administrator when attempting to use CA MICS StorageMate for the first time. Areas in which assistance may be needed include proper security level access (for example, RACF, Top Secret, ACF2), knowledge of the databases and data cycles that exist on your system, and knowledge about how the CA MICS Workstation Facility (MWF) is accessed at your site.

This section contains the following topics:

- [2.1 Security Authorization](#) (see page 23)
- [2.2 Database Information](#) (see page 24)
- [2.3 Levels of Summarization](#) (see page 26)
- [2.4 CA MICS Update Schedule](#) (see page 27)

2.1 Security Authorization

Users need the authorization to read both CA MICS 'sharedprefix' and 'prefix' data sets. These are the databases and files that constitute the CA MICS database. Because CA MICS StorageMate uses the facilities of both ISPF and MICF, you will need authority to use both of those facilities. Many facilities within CA MICS StorageMate support the use of data from an external data set, referred to as the Work Data Set. (This is a product-specific data set, not to be confused with a SAS Work Data Set.) You must have security authority to allow you to create and update this data set, although it is not required that the data set have any specific name structure. If you plan to include tape information from an external data set, you may want to verify that you have read access to this data set.

Most of these authorization needs should be satisfied by contacting your CA MICS System Administrator, your security administrator, and/or your TSO administrator.

2.2 Database Information

The CA MICS database is a repository of many types of information. For CA MICS StorageMate, the data of most interest includes RMF, SMF, VTOC, HSM, ASTEX, and catalog data. All of these data sources are kept at the DETAIL level, and also summarized at the DAYS, WEEKS, and MONTHS levels. Each of these levels, or 'timespans' as they are called in CA MICS, contains a different view of your data and can be used to answer different questions about your resource utilization. Because you may want to retain much of this data for several days or months, each of the files in each of the timespans may exist in a number of generations or cycles.

For example, if you decided to keep ten cycles of RMF data in your DETAIL timespan, these cycles would be numbered from 1 to 10. Cycle 1 would contain the data from the most current day, and Cycle 10 would contain data from ten days ago. The next time you run CA MICS, the new data would be stored in Cycle 1, existing Cycles 2-9 would bump up (or be 'aged') by one number, and Cycle 10 would be deleted. This cycle structure exists for every information area and in every timespan. The number of cycles maintained for any given file in any given timespan can vary, and is set by your CA MICS System Administrator. A thorough understanding of the structure and content of the database is central to the use of CA MICS StorageMate.

Files used by CA MICS StorageMate include a minimum of one cycle in the indicated timespan for each of the following CA MICS product files:

DETAIL

Application Extension for ASTEX*	-	ASTADS
MVS Batch and Operations Analyzer	-	BATPGM, BATJOB, BAT_JS
DASD Space Analyzer	-	VCADAA, VCA_VS, VCAVOA
HSM Inventory Analyzer Option*	-	HSMIMG, HSMVOL, HSBAC

DAYS

Application Extension for ASTEX*	-	ASTAVO
MVS Batch and Operations Analyzer	-	BATWDA
MVS Hardware and SCP Analyzer	-	HARDVA, SCPPGA
DASD Space Analyzer	-	VCADAA, VCA_VS
HSM Inventory Analyzer Option*	-	HSMIMG, HSBAC

MONTHS

DASD Space Analyzer	- VCADAA, VCA_VS, VCAVOA
HSM Inventory Analyzer Option*	- HSMMIG, HSBAC

* These files are only required if you want to use certain reports distributed with the StorageMate product. If these sources are not installed, you can still install and use CA MICS StorageMate, but some reports and report options will not be available.

Another important item of information you will need is the list of valid CA MICS database Unit IDs for your installation and a summary of the files and cycles contained in each. Make a note of this and keep it handy for future reference. You will need it to run the facilities within the CA MICS StorageMate product.

The CA MICS database complex can be segmented into different units to minimize operational problems, such as the coordination of incoming transmissions from remote data centers. Think of the complex as a partitioned data set (PDS), with each individual informational unit being similar to a member within the PDS. Files within these units might equate to records within each PDS member. Each of these unique units is identified in the CA MICS system by a one-character identifier, known as a Unit database ID (Unit DBID).

An example might be to have the following configuration of CA MICS products and unit databases:

CA MICS Information Area	Unit Database	Data Center Location
SMF, RMF, AST	D	Dallas
VCA, HSM	V	Dallas
SMF, RMF	N	New York
VCA, HSM, AST	Y	New York

These unit ID letters tell you where to find the data you will need to run your inquiries. For example, to run a report that requires only VCA data from the New York center, use DBID Y. To run the same report for Dallas, specify DBID V. To combine VCA data from both centers in the same report, specify both Y and V.

2.3 Levels of Summarization

If you are new to CA MICS, you may not know that it provides several views of your data and performs automatic summarization of the data at each higher view. The levels of information include DETAIL, DAYS, WEEKS, and MONTHS. The DETAIL level contains the unsummarized data, but in a format that is much easier to use than the original form. At this level you can get information about individual RMF intervals, individual jobsteps, and individual data sets. The DAYS level is the first level of summarization. It normally is summarized by hour. Therefore, if your RMF interval is 15 minutes, each hour record in the database will be representative of four original interval records, but will summarize the activity on your system for one hour.

Each file in each timespan can exist in several generations, or cycles. The number of cycles per file per timespan is defined by your CA MICS System Administrator. The number of cycles kept is usually influenced by the amount of storage space required by the file. Larger files normally have fewer cycles kept, to minimize the space requirements of CA MICS. Many of the files used by CA MICS StorageMate are usually small in size, so several cycles could be maintained with a minimum of impact.

You should obtain the normal processing schedule for your installation from your CA MICS System Administrator. You can then be assured that you are viewing and analyzing the data you want.

2.4 CA MICS Update Schedule

There are two basic methods of collecting storage data with the VCC scan job: (1) run it daily during off hours or, (2) run it periodically during the day to get finer granularity of information. For example, some installations combine the two methods, scanning temporary storage pools hourly and scanning permanent VSAM, DB2, and QSAM pools at midnight.

If your installation runs multiple scans in a day or processes multiple days within one cycle (for example, process Saturday and Sunday scans together), there is a customization option you should be aware of. This has to do with the way the product deals with repetitively collected data. For example, if you scan your temporary storage pool at the start of each of your three daily shifts, each cycle will contain three observations for each volume within this pool. Attempts to use this data with StorageMate reports will produce results showing three entries for each of these volumes. To avoid this problem, CA MICS StorageMate provides an option that can be set to eliminate this repetitively collected data. This option will either accept the latest time found in the input data and ignore all others, or it will use the latest time found for each individual volume or data set. The default option distributed with CA MICS StorageMate does not allow repetitively collected data to be checked. If you have repetitively collected data residing in each data cycle, as described above, you must make a change during the installation of the product to avoid having data repeated in your reports. Section 3.2 of this guide gives the details on how to make the change.

Chapter 3: CUSTOMIZING CA MICS StorageMate

Information on customizing the CA MICS StorageMate product is presented in the following subsections:

This section contains the following topics:

- [3.1 Defining Data Sets and Options](#) (see page 30)
- [3.2 Establish Global Processing Options](#) (see page 34)
- [3.3 Executing an Inquiry](#) (see page 39)
- [3.4 Volume Group Tables](#) (see page 46)
- [3.5 Data Set Group Tables](#) (see page 53)
- [3.6 Bypassing the Execution Run Time Screen](#) (see page 59)
- [3.7 Setting Up Your JCL Defaults](#) (see page 60)
- [3.8 Setting Up Your Graphic Defaults](#) (see page 61)
- [3.9 Data Filtering](#) (see page 63)
- [3.10 How to Enable Neugents technology in CA MICS](#) (see page 67)

3.1 Defining Data Sets and Options

There are two external data sets that need to be identified to CA MICS StorageMate by each individual user before they attempt to use the product. These data sets are the:

- o Work Data Set
- o Tape Data Set

A space reporting option also needs to be set, so StorageMate will know which method to use when reporting space usage for data sets managed by HSM.

The functions of these data sets and options, and the procedure for defining them to CA MICS StorageMate are described below.

The Work Data Set

ALLOCATING THE WORK DATA SET

Before using CA MICS StorageMate, you must allocate a data set that communicates with the StorageMate product. This data set, known as the Work Data Set, is a PDS that contains Volume Group Tables, Data Set Group Tables, generated JCL, data set name pattern masks, the sample ACS routines created from those masks, and other input or output members associated with other StorageMate facilities.

The Work Data Set is defined by choosing the "Data Sets and Processing Options" selection from the StorageMate and MICS Options menu, and then completing the "Work data set" area on the panel. The default Work Data Set name is "userpref.MICSSTOR.CNTL", where userpref is the default prefix for your TSO Userid. You can change this default name to any valid TSO data set name. If you specify a data set name that has not yet been created, you are shown a secondary allocation panel that creates the new data set for you. (This panel is similar to ISPF option 3.2). You should also establish the practice of periodically compressing the Work Data Set, and removing outdated or superseded members.

WORK DATA SET INTEGRITY

StorageMate users must recognize the importance of the Work Data Set, and know that it is not automatically backed up by CA MICS. If you do not have a regular backup procedure in place for these types of data sets, you may want to build

some backup jobstreams and have them executed on a regular basis either by the users or by a product administrator. While many types of products can be used to perform this backup, the example shown below uses the IEBCOPY utility to back up a Work Data Set to tape:

```
//BKUPJOB    JOB  (accounting-parameters)
//BKUPWDS    EXEC PGM=IEBCOPY
/*
/* PRIVATE WORK DATA SET BACKUP TO TAPE
/*
//SYSPRINT   DD SYSOUT=*
//SYSUT1     DD DISP=SHR,DSN=userpref.MICSSTOR.CNTL
//SYSUT2     DD DSN=userpref.MICS.MICSSTOR.CNTL(+1),
//           DISP=(NEW,CATLG,DELETE),UNIT=unit,
//           DCB=(userpref.MICSSTOR.CNTL),
//           LABEL=(1,RETPD=360)
//SYSUT3     DD UNIT=SYSDA,SPACE=(TRK,(3))
//SYSUT4     DD UNIT=SYSDA,SPACE=(TRK,(3))
//SYSIN      DD *
              COPY INDD=SYSUT1,OUTDD=SYSUT2
/*
```

The Tape Data Set

CREATING THE TAPE DATA SET

While no information about the tapes in your tape library currently exists in the CA MICS databases, a method has been provided for optionally including this information in some of the reports produced by the CA MICS StorageMate product.

You must first extract information about your tape library and place that information in a data set available to CA MICS StorageMate. Appendix C of this guide provides full details on the creation of this data set. While this data set may be prefixed with a TSO Userid, consider using a common prefix if multiple users will be running CA MICS StorageMate and accessing the data. Note also the data set will need to be refreshed regularly if you want it to contain the most current information about your tape library. You may want to schedule a daily production job that will do this, if you plan to use the data from this data set on a daily basis.

SPECIFYING THE NAME OF THE TAPE DATA SET

Once the external tape data set has been built, you must supply its name to CA MICS StorageMate. Choose the "Data

Sets and Processing Options" selection from the StorageMate and MICF options menu, and then complete the "Tape data set" area of the panel. The first time you invoke the option, the default name of the tape data set will be "userpref.MICSSTOR.EXTAPE", where userpref is the default prefix for your TSO Userid. Replace this name with the name you chose for your Tape Data Set. The name you specified will be saved in your profile and will be used whenever you specify that tape information be included in a StorageMate report. Even if your Tape Data Set is created using the default name shown above, you must still use this option screen to update your profile with the correct name.

Each individual user of CA MICS StorageMate must perform this step if they want to use tape data. Once your Tape Data Set is defined, you need not be concerned about this option again unless the name of the Tape Data Set changes, or you change your TSO Userid.

Specifying a tape data set name does not automatically include tape data in all future StorageMate reports. Whenever you execute a StorageMate report that supports input from multiple sources, you are given the option on the report screen to specify the report sources that should be used for that report execution.

HSM Space Reporting Options

OPTIONS AVAILABLE FOR HSM SPACE REPORTING

StorageMate supports reporting for data sets that are under control of HSM. Data sets fall under HSM control in one of two situations. In one case, data sets that have not been referenced recently have been moved from your online or "primary" volumes to HSM migration volumes. In the second case, one or more backup copies of the data sets on your primary volumes have been made on the HSM backup volumes. In either case, the space used on the HSM volumes for storing a data set is probably far less than the original amount of space allocated by the data set on the primary volume. This is because HSM uses such techniques as compression and data packing, and does not allocate any more space than is necessary to contain the data.

CA MICS maintains three different values to indicate the amount of space allocated by a data set under HSM control, and StorageMate may be set to use any of these options when reporting space allocated by data sets on HSM volumes. These

three values and their meanings are as follows:

ACTUAL - refers to the actual amount of space being used on the HSM volume to contain the data.

PRIMARY - refers to the original amount of space that was allocated on the primary volume by the data set before HSM took control of it.

ESTIMATE - indicates an estimate of the amount of space that would be required if HSM moved the data set back to a primary volume.

SELECTING THE REPORTING OPTION

Once you have selected one of the options described above, you must indicate your choice to CA MICS StorageMate. Choose the "Data Sets and Processing Options" selection from the StorageMate and MICF options menu, and then complete the "HSM space value" area of the panel with one of the three keywords shown above.

The option you have selected will then be used for all future HSM space reporting until you change it to a different value. When you run a StorageMate report that includes HSM data, a message in your MICSLLOG output report will remind you of the option you are currently using. Running the same report with different HSM space options is a good method for doing disaster recovery planning, or for measuring the effectiveness of HSM processing.

3.2 Establish Global Processing Options

The activities described in this section must be performed when the CA MICS StorageMate product is installed. They can be performed later, however, if you want to change any of the global processing options. Do this with caution, however, because these changes will apply to all users of the product.

If you are installing the StorageMate product for the first time, do not perform the activities described in this section until directed to do so by the checklist in Appendix B.

The global options described here are set by the person installing the product and will affect the way the product operates for all users. Each global option is set by modifying a member called \$STGEXIT that resides in the sharedprefix.MICS.SOURCE data set. Review each suggested change below to see if they apply to your installation. You may want to contact your CA MICS System Administrator if you have questions about them. Before making any of the changes suggested below, save an original copy of your \$STGEXIT member in your sharedprefix.MICS.LOCALMOD.CNTL data set.

Debugging Options

A statement within the \$STGEXIT member appears as follows:

```
%LET STGDEBUG = N;
```

This statement sets CA MICS StorageMate debug mode to active (= Y) or inactive (= N). Currently, the only function of debug mode is that it issues a warning message when it detects a data set name that is not found in a specified Data Set Group Table. You may consider activating debugging mode (= Y) when defining your initial Data Set Group Tables and then deactivating it (= N) later.

SAS Source Options

A variable named STGSASOP is defined within the \$STGEXIT member and is used to set the SAS source options for every CA MICS StorageMate report that is requested. As defined, the STGSASOP variable contains no option statements. If you want to permanently modify the SAS execution options for all CA MICS StorageMate reports, you can set the STGSASOP statement to include the options that will appear on the SAS OPTIONS statement. The word OPTIONS should not be coded, but

only a list of the options that will be activated. A common use of this facility would be to request the printing of source statements before calling the support group for a detected problem. An example of how you might do this is as follows:

```
%LET STGSASOP = SOURCE MACROGEN;
```

Use of Asterisk for Data Filtering

The CA MICS StorageMate Data Filtering panel allows the filtering of several data elements that contain character values, such as data set name or volume serial number. This panel will also allow generic specification of character values, also known as prefixing. The way the product is distributed, this generic specification is requested by using an asterisk as the last character specified. For example, you can select a report and ask for filtering by volume and specify the following values: SYSRES SP00L1 TSO* IMS*. These values would then select volumes with serial numbers SYSRES, SP00L1, plus any numbers with the prefix TSO or IMS. Specifying TSO without the asterisk would only select a volume if it had a serial number of TSO followed by three blanks. While this is not common in most installations, it is possible.

The use of an asterisk was chosen for consistency. Many other ISPF and storage management products use an asterisk as the last character to indicate a generic or prefix value.

An option is provided, however, that makes CA MICS StorageMate always treat every value as a prefix, eliminating the need to enter an asterisk as the last character. Find a statement in the \$STGEXIT member that appears similar to the following:

```
%LET STGASTER = Y;
```

As distributed, the value of 'Y' indicates that an asterisk must always be coded as the last character of a filtering value if it represents a prefix. Change this to 'N' if you want all short values to be considered prefixes.

Duplicate Data Options

The StorageMate product is distributed assuming that no repetitive data will exist in any of the data files used by

the product. Most installations run VCC once per day and convert that data into a VCA cycle. Some installations, however, have chosen to run VCC multiple times throughout the day and have all that data combined into one VCA daily run and one VCA/HSM cycle. An option is present in CA MICS StorageMate that can eliminate this repetitive data, but this option must be requested. Locate the statement within the \$STGEXIT member that appears similar to the following:

```
%LET STGREPOP = NONE;
```

Change the value following the equal sign (=) to be one of the following:

NONE

This option causes no checking to be done for repetitively collected data. Code this option if each VCA cycle represents only one execution of VCC for each system.

MERGE

Code this option to merge all data so that only the latest time stamp for each volume or data set will be included. For example, if you ran VCC at noon for volumes A and B and at 6 p.m. for volumes B and C, the noon data for volume A and the 6 p.m. data for volumes B and C would be included in the analysis. This is probably the best option to use if you have repetitively collected data.

LAST

Code this option if you want only data from the last time stamp found within the data to be used in the analysis. For example, if you ran VCC at noon for volumes A and B, and at 6 p.m. for volumes C and D, only data for volumes C and D would be included. Use this option if you run VCC throughout the day for selected volumes, and then at the end of the day for all volumes.

As distributed, this option is set to NONE, so that no checking for repetitively collected data is performed. Note, however, that there is one exception to this process. Some StorageMate reports allow you to enter a range of cycle numbers when you specify the report options. In this case, if you request multiple cycle numbers you would get repetitively collected data, even if each VCA cycle represents only one execution of VCC. To avoid this problem, CA MICS StorageMate will automatically invoke the MERGE

option if:

- 1) Option NONE was specified or defaulted in the \$STGEXIT member and
- 2) A range of cycle numbers was entered on the report options panel when executing a report.

If you code an option to eliminate duplicate data and no duplicates are found, there will be no effect on the final output. Note, however, that extra processing steps are added for this feature that will use more time and resources than would be used if STGREPOP was set to NONE.

Report Titles

The default title lines appearing at the top of each CA MICS StorageMate report are all defined in the \$STGEXIT option member. Thus, you can change these definitions if you want to have alternate titles on your reports. Use this option with caution, however, as it will apply to all reports executed by all users of the product. Note also that an alternate title may cause confusion when requesting technical support for a problem with one of your reports.

The default titles that appear on each report are in the following format:

- (title 1) Report Title (REPORT-ID)
- (title 2) CA MICS StorageMate Vx.x - Specific Report
- (title 3) Installation Name - Date

Most of these defaults can be changed by modifying control statements in the \$STGEXIT member. The statements shown below already exist in \$STGEXIT, and just need to be modified to contain the correct values:

- (#1) %LET STGExx01 = Report Title;
- (#2) %LET STGExxyy = Specific Report;
- (#3) %LET STGTTL02 = CA MICS StorageMate Vx.x;
- (#4) %LET STGTTL03 = Installation Name;
- (#5) %LET STGTTLDT = NO;

Line #1 exists for every separate StorageMate report, and specifies the value of the title 1 line. STGExx represents the unique Report ID (MICF ID) of each report.

Line #2 exists for those reports that support a number of different output reports, such as a SUMMARY report and a

DETAIL report. The value STGExx is as described above, while the yy portion will be a numeric value 02-99.

Line #3 specifies the default value for title 2. The default is the name of the StorageMate product, and its current version number.

Line #4 specifies the default value for title 3. The default is the name of your installation, as specified by your CA MICS Administrator when CA MICS was installed.

Line #5 specifies if the date should appear in the 3rd report title line. Specify a value of YES or NO. Depending on your SAS options, SAS may print a date, time, and page number in the upper corner of each report page. But if this option is not active, you may want to activate this option and print the date on the 3rd title line.

For color graphics, you always have the option of changing title lines 2 - 4 on the report option panel. You should probably not use these global options for color graphics changes, unless you want these changes to apply for all users, or you cannot get the titles you desire with the options provided.

Production Reporting Options

All StorageMate reports will execute in a MICS Production Reporting environment. This is done by setting certain defaults that will be in effect for each report that is selected for production reporting. There are some defaults, however, that could never be set in a way that would be satisfactory for all users. One example is if a report was needed to access a Work Data Set. For this reason, some defaults are set by options in the \$STGEXIT member.

If you are planning to use some of the StorageMate reports under the MICS Production Reporting facility, you may need to define some production data sets and define those data sets to the \$STGEXIT member. A description of these options is not given here. If you are planning to use MICS Production Reporting with this product, please edit the \$STGEXIT member and look for the comment relating to options for the MICS production reporting facility. Review the options under that heading and make any necessary changes. The function of each option will be described by the comments in the member.

3.3 Executing an Inquiry

This section is designed to guide you through your first use of CA MICS StorageMate. It should be followed by every new user of the product. It will lead you through one of the less complex reports provided by the product, so that you can start to get acquainted with the look and feel of the CA MICS StorageMate panels and facilities. For maximum benefit, you should follow the steps in this section while at your terminal.

All StorageMate reports include online tutorial panels that can optionally be displayed as you are requesting a report. Full details about the reports themselves are contained in this guide.

The StorageMate product uses the CA MICS Workstation Facility (MWF) to provide online capabilities for specifying operational parameters and selecting reports. If you are a first time user of MWF, refer to the Planning, Installation, Operation, and Maintenance Guide (PIOM), Section 3.3.6 (Setting Up the CA MICS ISPF Environment) and Section 4.14 (CA MICS Workstation Facility (MWF)).

The procedure for executing a sample inquiry:

- 1) Make sure the CA MICS StorageMate product has been installed successfully on your system. If you are unsure about this or if you get any unexpected errors or error messages during this procedure, contact your system administrator.
- 2) Select the Management Support Applications option from the MWF primary panel (shown in Figure 3-1) by typing the appropriate number in the Option area of the panel and then press the ENTER key. Contact your system administrator if you are unsure how to display this panel.

```
----- CA MICS Workstation Facility (MWF) -----
Option ==>
```

- 1 - Documentation Access (DOC)
Browse, print, and/or cross reference CA MICS documentation.
- 2 - CA MICS Information Center Facility (MICF)
Compose and/or replay CA MICS database inquiries.

```
3 - SAS With CA MICS Libraries (MSAS)
    Use interactive SAS with CA MICS libraries and macros.
4 - Management Support Applications (APPL)
    Accounting and Chargeback, Capacity Planner,
    MVS Performance Manager, StorageMate, Network Service Reporter,
    and MVS Model Generator.
5 - CA MICS Administrator Facility (MAF)
    Operational status and tracking, installation, modification,
    maintenance, and authorization.
T - Tutorial
X - Exit

. . . . .
```

Figure 3-1. CA MICS Workstation Facility Menu

3) A panel should now appear with the title "Management Support Applications" on the first line. Select "StorageMate" from this panel (shown in Figure 3-2) by typing the appropriate number in the Option area of the panel. Then press the ENTER key.

```
----- Management Support Applications -----
Option ==>

2 - Accounting and Chargeback
    Modify parameters, execute operational jobs, access Enterprise
    IS Financial reporting, and enter application files data.
3 - Capacity Planner
    Define and update capacity planning database files and perform
    forecasting analysis.
4 - MVS Performance Manager
    Analyze performance of MVS systems.
5 - StorageMate
    Report and analyze storage resources.
6 - Network Service Reporter Administration
    Maintain Network Service Reporter options and outage schedule
7 - MVS Model Generator
    Create analytic queuing models for OptiModel.

. . . . .
```

Figure 3-2. Management Support Application Menu

- 4) The primary StorageMate panel should now appear. This panel has the title "StorageMate" on the first line and is the first panel you will see when you enter the StorageMate product. From this panel, select option 0 to build your Work Data Set as discussed in Section 3.1.

```

+-----+
|----- StorageMate -----|
| Command ==>                |
| Please use Option 0 to define your Work Data Set |
| before continuing          |
|      0 - StorageMate and MICF Options             |
|      Review/Modify StorageMate and MICF Options |
|      and Defaults.                               |
| ==> 1 - Storage Administration                     |
|      Manage Problems, Data Sets, Volume Pools,  |
|      and Space Availability.                     |
|      2 - Performance Management                   |
|      Monitor Storage Performance.                |
|      3 - Capacity Management                     |
|      Analyze Storage Resource Usage.             |
|      4 - SMS Migration and Tracking               |
|      Plan, Implement, and Monitor System Managed |
|      Storage.                                     |
|      5 - Reports and Graphics                     |
|      Review Reports and Graphics Saved Previously. |
|      T - Tutorial                                |
|      X - Exit                                     |
+-----+

```

Figure 3-3. StorageMate Primary Panel

- 5) The next step is to select the report area you want to use. For this example, enter the number 1 in the Command area of the StorageMate Primary Panel and press ENTER. This tells CA MICS StorageMate that you want to select Storage Administration functions.
- 6) A panel should now appear with the title "Storage Administration" as the first title line. This is the main panel from which all functions related to Storage Administration can be selected. For this example, enter the number 1 in the Command area and press ENTER. This tells CA MICS StorageMate that you want to select Problem Management. If you have not defined your Work Data Set, the panel will display a message telling you to use Option 0 to do so before continuing.

- 7) A panel should now appear with the title "Problem Management" as the first title line. For this example, enter the number 2 in the Command area and press ENTER. This tells CA MICS StorageMate that you want to select Volume Problems.
- 8) Depending on how your options have been set, you may see a panel at this point with the title "MICF Execution." This represents the options that MICF will use to execute the Volume Problems report you have just selected. Section 3.6 explains how to bypass this panel once you have set the options to meet your specifications. If this panel is displayed, the most important item to complete is the Inquiry Execution Mode section. Move the cursor to this area and enter FOREGROUND (or just the letter F). Then press ENTER.
- 9) A panel now appears with the title "Volume Problems." This is the StorageMate panel used to indicate the options that should be in effect for this particular execution of the report. Each report produced by CA MICS StorageMate has a panel similar to this one. Running the same report with different execution options will often give you many different views of your data and can be used to solve many different types of problems.
- 10) Type the word HELP in the Command area of the panel and press ENTER (or press the HELP Program Function Key). This displays a series of tutorial panels that explain the function of this report and explain the rules for completing the execution options panel. Each CA MICS StorageMate report comes complete with a set of tutorial panels similar to these to assist you in running the report and understanding its function. After using the tutorial feature, press the END Program Function Key (or type END in the Command area and press ENTER) until you are back at the Volume Problems execution options panel.
- 11) The next step is to specify the options that should be used for producing the Volume Problems report. These options are specified on the execution options panel currently being displayed. Each StorageMate report has an execution options panel similar to this one. The ability to change the options before executing the report allows you to run the same report several times with CA MICS Input Definition section of the panel identifies the CA MICS data that should be read as input to the report. Specify a question mark (?) in the Unit DBID(s)

section of this panel and a numeric value of 01 in the Cycle section. The Report Options section of the panel allows changes that affect the output that is produced by the report and the appearance of that output. This report has only one option and that is to produce a jobstream that could be used to correct certain problems discovered by the reporting program. For this particular example, leave all the selection areas in the Report Options section blank.

- 12) In this example, you placed a question mark (?) in the Unit DBID(s) area because you did not know the units that are valid for your installation. If only one unit is valid, the one-character DBID value will be placed in the Unit DBID(s) area and the Volume Problems execution panel will be displayed again. You must then use the END key again to start report execution. If more than one unit has been defined, a panel with the title "Data Base Selection" will be displayed. All of the units that are valid for your installation will be displayed on the panel. Move the cursor down the left side of the panel and place a S next to the unit containing the VCA data you want to use as input to this report. Then press the ENTER key.
- 13) Notice the key legend that appears as the last line of the panel. It lists and briefly describes the functions of the keys that are valid for this panel. Notice that the DOWN key is valid for a function called Filter. This function allows you to filter, or subset, the input data that will be used for the report. For example, you may want to exclude data for volumes with a serial number prefixed with "SYS", or include data sets created within the last 30 days. Use the DOWN key to display the StorageMate Data Filtering panel. This panel is common to all reports, although the type of filtering varies for each report. Notice that the panel has three columns; all will accept up to 18 rows of filtering statements. Each statement consists of an Element to be tested, the comparison Operation to be done, and the Values against which the element will be tested. No actual data will be filtered for this example, but you may want to use the HELP function to learn more about the filtering options available. When you are finished with the filtering panel, follow the key legend at the bottom of the panel, and use the UP key to return to the main report panel.
- 14) When you have completed the execution options panel, you can either press ENTER or the END Program Function Key.

If you press ENTER, the values you entered will be validated but report execution will not begin. If you use the END key, the values you entered will be validated, and then report execution will begin if all values are valid. You may want to use ENTER to make sure the values you entered are correct, followed by END to start the report. Or you may just want to press END directly. For this example, press the END Program Function Key.

- 15) When report execution begins, the Volume Problems panel will disappear and several status panels will be displayed. These panels contain a large comment block in the middle of the screen describing what is being done. A message in the upper right-hand corner of the panel also gives additional information.
- 16) When you select FOREGROUND execution, there is a slight delay while allocations are performed, SAS is invoked, and the report is produced. After a short time, the report should appear at your terminal. You can use the ISPF UP, DOWN, LEFT, and RIGHT keys (or commands) to browse through the report. When you have finished examining the report, you should keep pressing the END key to get you back one level, until you are all the way back to the Workstation Facility (MWF) main panel.

You have just completed your first StorageMate report. You can now expand on the knowledge just learned to request other reports and try other options. If questions occur during this process, remember to use the tutorial facility or the report descriptions in this guide. If you had problems generating your report, or it did not operate as expected, contact your system administrator. You may have specified an option incorrectly or your MICF options may not be set correctly. Your system administrator should be able to help you with either of these problems.

Accessing StorageMate Reports through MICF

The CA MICS StorageMate product is designed to guide you through normal storage administration functions with a group of selection menus organized by storage function. This provides a results-oriented approach for those who may be new to the StorageMate product, or to storage administration in general. If you are a more experienced StorageMate user, or if you know the specific StorageMate report you want to run, an alternate method of accessing reports is provided through

MICF. In order to do this, you must select the MICF option from the MWF (CA MICS Workstation Facility) menu. After you enter the MICF facility, select the Data Base Inquiries option. Display the inquiries in the STORAGE Catalog Group of the Shared catalog. All the StorageMate reports should now be listed, and you may select and execute the report or reports you desire. If you have questions about using this procedure, contact your system administrator.

3.4 Volume Group Tables

Your duties have grown more complex as the number of DASD volumes installed has increased. While management on an individual volume basis is possible with a few dozen volumes, it becomes more complex when the number of volumes grows to several hundred or more. This management task is also complicated by the fact that several different types of devices may be installed at any one time.

One concept that has evolved to assist in the complex task of managing many volumes is the volume group. A volume group represents a series of volumes with a common function. Examples of common volume groups might be as follows:

- o MVS volumes used for paging
- o All dedicated volumes belonging to payroll
- o Volumes used for new TSO data sets
- o Volumes used for temporary work space

Several of the facilities within CA MICS StorageMate allow you to report at the volume group level. This makes reporting more effective, because it is based on the LOGICAL function being performed by each group, rather than by the PHYSICAL identification of a single volume.

The vehicle used to communicate your volume groups to CA MICS StorageMate is a Volume Group Table. This table is always required if you want to use reports that do volume grouping and you are in a non-SMS environment. For SMS users, each volume controlled by SMS is assigned to a Storage Group. StorageMate will use any non-blank Storage Group name found in the data as the volume group name. Thus, SMS users only require a Volume Group Table to assign groups to those volumes not managed by SMS. The one exception to this is the SMS user that desires to simulate the effect of Storage Group changes, as a Volume Group Table can also be used to perform this function.

In summary, users desiring to use volume grouping reports will need to define Volume Group Tables under three conditions:

- 1) No volumes are managed by SMS.
- 2) Some volumes are managed by SMS, but grouping is also desired for the volumes that are not managed.
- 3) Volumes are managed by SMS, but the user wants to ignore

the SMS Storage Group and group all volumes by means of a Volume Group Table.

After you have defined your Work Data Set, you can use options from the CA MICS StorageMate and MICF options menu to create and maintain your Volume Group Tables.

Even though most users will only need to define one Volume Group Table, you can create as many as you need. You may want to have separate tables for multiple sites or to have multiple tables that give different views of the same volumes at one site. Note, however, that each execution of a CA MICS StorageMate report can only access one Volume Group Table.

As already stated, installations that have fully implemented SMS may not need to define Volume Group Tables. This is because the SMS Storage Group assigned to each volume is used by CA MICS StorageMate as the Volume Group. Unless you are using a Volume Group Table designed for simulation, CA MICS StorageMate only uses the table if the volume being processed has no Storage Group name assigned to it (that is, the STORNAME data element value contains blanks). This means that users of SMS or any other system that supplies a value to STORNAME may not need to define Volume Group Tables.

SMS users may want to use Volume Group Tables under the following conditions:

- o Even installations with full SMS implementation will probably have some volumes that are not controlled by SMS, such as system volumes or JES volumes. You may want to define a Volume Group Table just to assign Volume Groups to these volumes. When creating this type of Volume Group Table, specify on the options panel used to create the table that SMS-managed volumes should NOT be included. The use of this option will cause only those volumes with no SMS Storage Group assigned to be included in the Volume Group Table that is generated. It will then only be a simple matter to update the Volume Group Table and assign group names to the few volumes that are not controlled by SMS.
- o SMS installations can use Volume Group Tables to simulate alternate Storage Group configurations. Users would first generate a Volume Group Table using the full 6-character Volume Serial Number, and specifying that SMS volumes SHOULD be included. This will create a Volume Group Table with all volumes represented, and with the SMS Storage Group name for each volume presented as the Volume Group.

It will also cause the table to be searched for all volumes, including SMS-managed ones. You should then edit the Volume Group Table and make the necessary changes to Storage Group assignments. After this modification, you may want to run multiple StorageMate reports, changing the table definitions until you achieve the desired results. The ability to update the Volume Group Table allows you to simulate many different combinations of Storage Groups before making any actual changes to your volume groupings.

Please note again the effect of the SMS volume option when using CA MICS StorageMate to build a Volume Group Table. If you specify 'N' (No), only volumes not managed by SMS will be included in the table, and use of that table will cause the table to be searched for unmanaged volumes only. If you specify 'Y' (Yes), all volumes will be included in the table, and use of the table will always cause the table to be searched for all volumes, including managed ones.

Installations that are migrating to SMS can use CA MICS StorageMate to simulate Storage Groups before they are actually implemented. This simulation enables you to use the StorageMate product to assist in the SMS migration process.

Example of a Volume Group Table

This installation not yet implemented SMS, but they have established the following standards for their volumes:

- o The first three characters of the volume serial number describe the usage of the volume.
- o The prefix HSM is assigned to HSM Level 1 volumes.
- o The prefixes MVS and MXA are assigned to operating system volumes.
- o The prefix TSU is assigned to TSO storage volumes.
- o The prefix Snn is assigned to all IMS production volumes, where nn represents a numeric value.

This installation generates a Volume Group Table, specifying that the first three characters of the volume serial number be used. Because most of the volumes found will be IMS volumes, the default volume group assignment is specified as IMSPROD. The generated table appears below:

VOL=HSM*	IMSPROD
VOL=MVS*	IMSPROD
VOL=MXA*	IMSPROD
VOL=S08*	IMSPROD
VOL=S24*	IMSPROD
VOL=S28*	IMSPROD
VOL=S32*	IMSPROD
VOL=TSU*	IMSPROD

The installation will now update the Volume Group Table to contain the correct group assignments. No IMS volumes must be updated, because they already have the default assignment of IMSPROD, which is correct for those volumes. Other volumes should be assigned to group names that reflect their function. The updated table appears below:

VOL=HSM*	HSMLV1
VOL=MVS*	SYSTEM
VOL=MXA*	SYSTEM
VOL=S08*	IMSPROD
VOL=S24*	IMSPROD
VOL=S28*	IMSPROD
VOL=S32*	IMSPROD
VOL=TSU*	TSODS

Any CA MICS StorageMate report that uses this Volume Group Table will do grouping using the volume group names shown in the right-most column of the table. While the above table will work correctly as shown, additional changes could be made to reduce the size of the table by using even shorter prefixes for some groups. The table shown below will function correctly, but is smaller than the previous table:

VOL=HSM*	HSMLV1
VOL=M*	SYSTEM
VOL=S*	IMSPROD
VOL=TSU*	TSODS

Detailed information about Volume Group Tables is available in the following subsections:

- 1 - Creating a Volume Group Table
- 2 - Modifying a Volume Group Table

3.4.1 Creating a Volume Group Table

Volume Group Tables are created as follows: Select the Create Volume Group Table option from the CA MICS StorageMate and MICF Options menu. Selecting this option executes a CA MICS StorageMate facility that creates your Volume Group Table. Depending on your environment, the next panel you see may have the title "MICF Execution". If this panel appears, you should complete it and press the ENTER key again. The Create Volume Group Table report options screen should then appear.

This screen has three required fields. The first is the Unit DBID(s) field. In this field, enter the CA MICS unit database ID (DBID) discussed in Section 2.2. You can also enter a question mark to see a list of the valid DBIDs already defined by the CA MICS System Administrator. The second required field is the Cycle field. Enter the number of the cycle that will be read to create your table. The third required field is the Volume Group Table name field. Enter a meaningful name for the table you are about to create, because you will need to enter this name again when you request a CA MICS StorageMate facility that does volume grouping.

The Create Volume Group Table screen also has some optional fields. In the Default description field, you can overtype the previous or default Volume Group name and replace it with a name that will represent the pool with the largest number of volumes assigned to it. This name will appear as the Volume Group name assigned to each volume in the table you are creating. You must then edit the Volume Group Table to assign more meaningful names. So if you specify the name of the largest group as the default name, it will minimize the amount of editing you will have to do later. For example, if 90% of your volumes are IMS production volumes, specify a default name such as IMSPROD. Then you will only have to change 10% of the volume assignments when editing.

Another optional field deals with the portion of the Volume Serial Number (VOLSER) to use in creating the table. The default value is 6, which will include the entire VOLSER value and will also use the SMS Storage Group name instead of the default name you specified, if a valid Storage Group name is present. If you want to use less than six characters, overtype the default with the value corresponding to the number of characters you want to use. For example, if you have a SORTWORK pool that consists of volumes SORT01, SORT02, SORT03, you can select all three volumes by changing the

value to 4, which causes the table to display an entry of SORT*, which will select all volumes prefixed with SORT.

The last optional field determines whether volumes already assigned an SMS Storage Group should be included in the table. The value you specify will also affect the way the table is searched when it is used within a report, as described in Section 3.4 of this guide. Unless you are simulating changes to your Storage Group definitions, you will probably wish to code 'N' (No) for this option.

As with all StorageMate report option panels, you can use the ISPF UP and DOWN keys to display the StorageMate Data Filtering panel. When creating a Volume Group table, data filtering can be used to limit the volumes or systems that will be included in the resulting table.

After you have modified the Create Volume Table report options screen, press ENTER to validate your changes. If no error messages appear on your screen, all the values you have entered are correct and you are ready to create your table. To do this, press the END key, or type END on the command line of the screen and press ENTER.

3.4.2 Modifying a Volume Group Table

After creating a volume group table, you must then modify it to make all the Volume Group names accurately reflect the groupings. To modify the table, select the Work Data Set Maintenance option from the StorageMate options menu.

The system displays a screen that prompts you to select a table name. The name of the last Volume Group Table you created or used should appear elsewhere on this screen for your information. You can enter the name of the table you want to edit, or leave it blank to get a list of all the members within your Work Data Set. From that list you can then select the table name to be updated.

When you have selected a table to update, CA MICS StorageMate places you in ISPF edit mode. You can modify the table as in any ISPF edit session. Overtyping the default pool name for those volumes in which the default Volume Group name does not apply, replacing the default with the proper Volume Group. If you have a large number of volumes to be grouped and they all start with the same prefix, place an asterisk (*) after the prefix on the first one and then delete the other entries with that same prefix. You will then have identified all the volumes of that prefix with the associated Volume Group.

There are no right or wrong names when assigning Volume Groups. These names have no relationship with any existing MVS or SMS names (with the exception of Storage Group), and are simply used to make Volume Group reporting possible. The Volume Group names you assign should be meaningful to you, but need not match another name within your environment. When your updates to the Volume Group Table are complete, use the END PF key or the END command until you return to the CA MICS StorageMate options menu.

3.5 Data Set Group Tables

While many reports within the StorageMate product deal with data sets on an individual basis, others summarize the use of your storage by groups of data sets. While individual data set reporting is valuable for detecting specific performance or usage problems, grouping is more valuable for detecting trends. In order for these data set grouping reports to provide you with usable information, the data set grouping must accurately reflect how your storage is used by each group.

When you request a CA MICS StorageMate report that does data set grouping, the report options screen allows you to specify that the grouping be done in one of three ways:

- 1) Grouping is to be done using combinations of the VCA account codes (VCAACT1 - VCAACT9) that exist within your data. Each day, as part of VCA daily processing, an exit is called for each data set processed. This exit, which is written by your installation, examines various characteristics about the data set and then sets one or more account fields that are retained with the data. One of these codes might identify the division, while a second identifies the department, and a third identifies the individual within the department.

Of the three summarization methods, this one is probably the easiest to maintain, and provides the most useful information. You may want to review or modify the current structure or content of your account codes to make them more meaningful for StorageMate reporting.

- 2) Another way to group data sets is by using one or more of the nodes that compose each data set name. This grouping method is valuable if you have accurate data set naming standards in place. For example, if all data sets have a first node identifying the department or group that owns them, such as PAYROLL or SALES, you will have an accurate grouping system already in place by using the first node of the data set name.
- 3) The final data set grouping method supported by CA MICS StorageMate is the Data Set Group Table. This table allows you to associate meaningful labels with different combinations of data set name prefixes. For example, PAY.SITE1.* data sets could be labeled "Miami-Payroll" and PAY.SITE2.* data sets could be labeled "New-York-Payroll". Facilities accessed through the

CA MICS StorageMate and MICF options menu assist you in creating and maintaining Data Set Group Tables.

Choosing the method of data set grouping is not a global decision. You can choose from the three methods described above each time you execute a StorageMate report. You can also alternate between methods, depending on the situation. For example, you can run a report using the first accounting field and detect a problem with the data sets in the payroll department. You might then run the report a second time, selecting only payroll data sets and using a Data Set Group Table to further identify each payroll data set at the owner level.

After you have defined your Work Data Set, you can use options from the CA MICS StorageMate options menu to create and maintain your Data Set Group Tables.

As with Volume Group Tables, you can create as many Data Set Group Tables as desired, either to represent data from multiple sites or to provide different levels of detail for the same data. For example, one table can assign ownership at a very high level, such as division, while others could show ownership at lower levels, even down to the individual user.

Detailed information about Data Set Group Tables is available in the following subsections:

- 1 - Creating a Data Set Group Table
- 2 - Modifying a Data Set Group Table

3.5.1 Creating a Data Set Ownership Table

To create a Data Set Group Table, select the "Create Data Set Group Table" option from the StorageMate and MICF Options menu.

You can, optionally, see the MICF Execution panel described in Section 3.6 of this guide, after which you should see the Create Data Set Group Table report options screen. It is similar to the screen you saw when creating a Volume Group Table. You must specify the Unit DBID and Cycle number identifying the data to be used. Another option in the CA MICS Input Definition area of the panel allows you to identify the sources of data to be used when building the table. In this area you must specify one or more letters, each separated by at least one blank space. Place a question mark (?) in this field and press ENTER to get a list of the valid values for this particular function. As with the creation of a Volume Group Table, you can use the StorageMate Data Filtering panel to subset the input data to be used for creating your table. You may wish to include only those data sets residing on certain volumes or systems, or only those data sets having a specific data set name structure.

As with the creation of the Volume Group Table, you must specify the name of the table that will be created, as well as the 16-character default description that will appear on each entry. Later, you will change this description to contain a more meaningful value for each prefix.

Another field specific to the creation of data set tables is the 'Number of DSNodes to use' field. In this field, enter a numeric value from 1-6 that corresponds to the number of nodes that should be considered for each unique table entry. Use any value greater than 1 with some caution, as this could create a very large table.

The final option is whether mask characters will be supported in the table that will be created. If mask characters are not allowed, each entry in the table consists of a data set name prefix. If mask characters are allowed, the table is initially created as a table of prefixes, but can be changed to contain any valid mask character. An advantage of mask characters is that the tables tend to be smaller and easier to maintain. A disadvantage of mask characters is that they take longer to process when running a report that uses a table containing them. This may cause reports to run longer and consume more processor time. The decision concerning mask characters need not be universal, as you may create and

use both types of tables.

When you have completed your entries on the Create Data Set Group Table screen, press ENTER to validate your changes; then use the END PF key to actually create the table.

3.5.2 Modifying a Data Set Ownership Table

After the Data Set Group Table is created, you must assign meaningful group names to each data set prefix. Use the "Work Data Set Maintenance" option from the CA MICS StorageMate and MICF Options menu to edit the table you have just created. You will see an ISPF edit screen showing two columns. The left column contains the data set prefixes that exist on your system, and the right column contains the default group description you assigned to each group.

You may want to edit the left column to reduce the number of entries in the table. Even though the table is created with full data set nodes, the trailing asterisk that is present for each entry can follow any character. For example, assume you specified that each table entry should be based on the first two nodes of the data set name, and the resulting table contained the following seven entries for payroll data sets:

DSN=PAY.CNTLBOB.*	DEFAULT-GROUP
DSN=PAY.CNTLSAM.*	DEFAULT-GROUP
DSN=PAY.MASTNEW.*	DEFAULT-GROUP
DSN=PAY.MASTOLD.*	DEFAULT-GROUP
DSN=PAY.QTEST.*	DEFAULT-GROUP
DSN=PAY.RTEST2.*	DEFAULT-GROUP
DSN=PAY.ZTEST.*	DEFAULT-GROUP

The first two entries represent data sets belonging to Bob and Sam, the two production control people that run payroll. The next two entries represent production data sets. The final three entries represent data used in payroll testing. After editing, these seven entries might be reduced to the following three:

DSN=PAY.CNTL*	PAYROLL-CONTROL
DSN=PAY.MAST*	PAYROLL-PROD
DSN=PAY.*	PAYROLL-TEST

When using the table to determine data set grouping, the entries are tested in the order they appear in the table, and the first match stops the table search. Thus, in the above example, all payroll data sets that do not start with a payroll control or production node will be assigned to the payroll test group.

Use the right column of each table entry to describe the data set group. This description MUST NOT contain imbedded blanks--use a dash or slash as a separator. If you use a space as a separator, CA MICS StorageMate will stop reading

your description at the space, and will ignore the remainder. The description may be up to 40 characters in length, but some reports truncate after the first 16 characters. Make sure the first 16 characters are unique for each group.

The example above assumes you are using a table that does not support masking characters (this decision is made when the table is created). The example shown below is from a table that does support mask characters:

DSN=OPER1.**	DEFAULT-GROUP
DSN=OPER2.**	DEFAULT-GROUP
DSN=OPER3.**	DEFAULT-GROUP
DSN=TESTATM.**	DEFAULT-GROUP
DSN=TESTPAYR.**	DEFAULT-GROUP
DSN=TESTS1.**	DEFAULT-GROUP

Notice that the syntax is different for a table containing mask characters, and that a prefix is indicated by two trailing asterisks. The modified version of this table is shown below:

DSN=**,ISP%LIB	ISPF-APPL-DEVELOPMENT
DSN=**,*SYSP*,**	SYSTEMS-PROGRAMMING
DSN=OPER%.**	SYSTEM-OPERATORS
DSN=TEST*,**	PRODUCT-TESTING

Two new entries have been inserted at the beginning of the table, one to group ISPF libraries and one to group data sets belonging to systems programming. The original six entries for operations and product testing have been reduced to two entries.

The masking characters that are supported are further described in the Facility Procedures section under the heading "Instructions for Creating and Maintaining the Pattern File" at the end of Chapter 4 of this guide.

You have just created a Data Set Group Table. This table now exists as a member within your Work Data Set and can be used to assign data set ownership in any StorageMate report that supports this facility.

3.6 Bypassing the Execution Run Time Screen

As described previously, a "MICF Execution" panel may be displayed after you have selected a report but before the report options panel is displayed. The MICF Execution panel allows you to designate the execution mode of the inquiry. There are three execution modes: batch, foreground, and SAS DMS. The execution mode determines whether CA MICS StorageMate reports will execute while you wait and then display the results at your terminal, or whether a jobstream will be created that will produce the report in batch. The advantage of the first method (foreground) is that you can see your results after a short delay. The advantage of the second method (batch) is that you can continue working at your terminal while the report is produced, and the report will be printed rather than displayed on your terminal. If you select batch execution, the MICF Execution panel also allows you to specify the job card information to be used.

If you always run your reports in the same mode, or if you have a great number of reports to generate, you can set the mode defaults to the desired settings and then bypass the display of the MICF Execution panel. This is accomplished by changing the 'Execution Mode' default parameters in MICF. This default is 'Yes' for all new users, meaning that the MICF Execution panel will always be displayed.

The specific steps for turning off the panel display are:

- 1) From the main StorageMate menu, select option 0 to enter the options facility.
- 2) Select option 0, 'MICF Options', which allows you to modify the options MICF uses during inquiry execution.
- 3) Select the option 'User Profile Parameters' from the MICF Options panel.
- 4) Scroll down in the display until you reach the 'Execution Mode' selections. Select the mode (Batch or Foreground) and also change the 'Yes' to 'No' to suppress the display of the panel.
- 5) Scroll through other items on this menu to ensure other options are set correctly for your environment. For example, if you have selected batch execution of reports, verify the default job card information is correct.
- 6) Press your END PF key until you return to the CA MICS

StorageMate main menu.

The display of the execution mode panel can be turned back on by repeating this process, but changing the 'No' in step 4 to 'Yes'.

3.7 Setting Up Your JCL Defaults

Section 3.6 describes a technique for changing MICF defaults to suppress the MICF Execution screen when running a CA MICS StorageMate report. This involved using the MICF options facility to update the "User Profile Parameters". Many of these parameters also affect the batch JCL (job control language) that is generated when you run a report in batch mode. If you plan to regularly run StorageMate reports in batch mode, we suggest you review these options and make any necessary changes for your environment. The procedure to do this is as follows:

- 1) From the main StorageMate menu, select option 0 to enter the options facility.
- 2) Select option 0, 'MICF Options', which allows you to modify the options MICF uses during inquiry execution.
- 3) Select the option 'User Profile Parameters' from the MICF Options panel.
- 4) Scroll through the items on this menu to make sure all options are set correctly for your environment. Make sure that the default job card information is correct.
- 5) Press your END PF key until you return to the CA MICS StorageMate main menu.

3.8 Setting Up Your Graphic Defaults

GRAPHIC DEFAULTS

Many StorageMate analyses will produce either tabular or graphic results. Graphic output can be produced on either a color device or a monochrome printer. When you run a StorageMate report for the first time, the default color graphic options to be used will be set from your MICF Color Graphics options. Even though an override panel is provided by each StorageMate report that produces graphics, having to override a large number of options will create more work for you and more possibility of error. For this reason, you should review the MICF options relating to color graphics if you plan to generate color graphics from StorageMate reports. The procedure to do this is as follows:

- 1) From the main StorageMate menu, select option 0 to enter the options facility.
- 2) Select option 0, 'MICF Options', which allows you to modify the options MICF uses during inquiry execution.
- 3) Select the option 'Color Graphics Format Parameters' from the MICF Options panel.
- 4) You will be shown a list of the various color graphics devices defined to MICF. Place an 'S' to the left of the device you want to modify, and press ENTER.
- 5) Scroll through the items on this menu to make sure all options are set correctly for your environment. If you are unsure what certain options do, you may want to leave them unchanged until you have had more experience with StorageMate color graphics reports.
- 6) Press the END PF key until you return to the CA MICS StorageMate main menu.

EXECUTION DEFAULTS

Most inquiries in CA MICS StorageMate are most efficiently run in BATCH mode. The MICF user profile Execution Mode option specifies the execution mode to be used. Review the default execution mode as set by the CA MICS System Administrator. If it reads FOREGROUND, change it to BATCH for the duration of the StorageMate session. Be sure to change the execution mode back to FOREGROUND, if necessary, after your CA MICS StorageMate session is completed. The

steps to follow are as follows:

- 1) From the main StorageMate menu, select option 0 to enter the options facility.
- 2) Select option 0, 'MICF Options', which allows you to modify the options MICF uses during inquiry execution.
- 3) Select the option 'User Profile Parameters' from the MICF Options panel.
- 4) You will be shown a list of the user profile parameter titles. Look first at the 'Execution Mode Display' parameter. If it reads YES, you will be able to specify BATCH processing at the time of StorageMate inquiry execution. If it reads NO, locate the 'Execution Mode' parameter.
- 5) To specify BATCH processing as a default, enter BATCH at the prompt under the 'Execution Mode' parameter.
- 6) Press your END PF key until you return to the CA MICS StorageMate main menu.

Note that if your options are not set correctly for your graphic device, SASGRAPH may not be able to produce your graph. If this occurs, refer to the error message produced by SASGRAPH in the SAS log to diagnose the error. You may need to alter the number of horizontal or vertical axis points defined, alter the size of chart bars, or make other changes to your graphic options to be compatible with your output device. You may want to make these modifications temporarily on the report execution panel, and then make permanent changes in MICF when you are satisfied with the results.

3.9 Data Filtering

Reports are often more meaningful if you can limit them to just a subset of your input data. All StorageMate reports support the use of a common filtering panel, identified with the title "StorageMate Data Filtering." This panel may be accessed from each report options screen, using the ISPF UP and DOWN keys. This filtering panel allows you to subset, or filter, the input data to contain only those items of interest to you.

Use filtering with caution, as it is possible to produce reports that would be misleading, or to specify a set of filtering criteria that cause no data to be selected at all. If you have problems running a certain report, first try removing all filtering and then try again.

The filtering panel contains up to 18 filtering statements. Each statement consists of three parts; an Element, an Operator, and one or more testing Values.

The Element column is used to indicate the data item that is to be tested for filtering. Examples include the data set name, the volume serial number, or the DFSMS storage class.

The Operator column indicates the type of comparison that will be made. For character elements, the following types of operators are valid:

- EQ - Element must be equal to one of the values
- NE - Element must not be equal to any of the values

In addition to the operators shown above, the following operators are valid for numeric elements:

- LT - Element must be less than the value
- LE - Element must be less than or equal to the value
- GT - Element must be greater than the value
- GE - Element must be greater than or equal to the value
- RA - Element must be within the lower/upper range

The Values portion of each filtering statement should contain zero or more values against which the element will be compared. Multiple values are valid for EQ and NE comparisons. Only one value is allowed for the other operands, except for RA, which must have a lower and an upper limit specified. Unless your installation has changed the StorageMate default, character operands may end in a trailing asterisk (*) to indicate a prefix, generic, or "wild card"

value. The one exception to this rule is when testing the data set name, as that element supports values containing any of the valid data set name mask characters. With character elements, a special keyword of "NULL" is supported to test for a value of blanks or spaces.

For those StorageMate users familiar with the SAS programming language, the Data Filtering panel will also allow free-form SAS statements to be used in the filtering process. Enter an Element name of "SAS", and Operator of "EQ", and then the SAS statement to be included. The SAS Element is the only element that may appear more than once on the data filtering screen. You can use some or all 18 lines of the screen for SAS statements. You may also continue SAS statements, provided you don't mix them with other filtering statements:

Element	Oper	Value(s)	Element	Oper	Value(s)
-----			-----		
SYSID	EQ	TST* PROD	SAS	EQ	IF RECFM EQ 'U'
SAS	EQ	IF RECFM EQ 'U'	SYSID	EQ	TST* PROD
SAS	EQ	THEN DELETE;	SAS	EQ	THEN DELETE;
(CORRECT)			(INCORRECT)		

Additional examples of using SAS for filtering are shown later in this section.

An example will now be presented of a typical completed filtering panel:

```
Element Oper Value(s)
-----
SYSID    EQ  TST* PROD
SAS      EQ  IF SYSID EQ 'TSTZ' THEN SYSID = 'TEST';
DSNAME   NE  SYS1.** *.LIST
STORCLAS EQ  NULL
REFDATE  RA  01JAN91 31DEC91
ALLOCMB  GE  10
```

The first line selects data from the PROD System Identifier, plus any SYSID starting with "TST". The second line is a SAS statement that will change any SYSID with a value of "TSTZ" to a new value of "TEST". The third line will drop any data sets with the prefix SYS1 or the suffix LIST. The fourth line selects only those data sets with no SMS storage class assigned. The fifth line selects any data set referenced any time in 1991. The sixth line selects all data sets that occupy 10 or more megabytes of space. Note that each input record read for the report shown must meet all of the

criteria shown to be selected.

More Examples of SAS Filtering

As explained above, the StorageMate Data Filtering panel allows you to enter multiple lines of SAS programming language code. This makes the filtering and reporting capability of CA MICS StorageMate much more powerful. Any statements you code on the panel are placed into the data step for each data step used by the report. For example, if the report you execute uses the VCAVOA and HARDVA files, any SAS code you specify will execute for both the specified files. If you want your code to apply only to some of the files processed, a variable named "FILE" can be tested to determine the file being processed. This will be set to the last three characters of the file name being processed. For example, you could use the following code to filter within the VCAVOA file only:

```
IF FILE EQ 'VOA' AND VOLSER EQ 'SYSRES' THEN DELETE;
```

While any SAS statement can be included with this facility, remember that the code you specify will be included as part of one or more data steps. In most cases, this code will perform additional filtering not possible with the filtering options present on the panel. Do NOT include SAS code that will cause the current data set to be terminated, such as the inclusion of a RUN statement or a SAS PROC statement.

Examples of valid SAS code are shown below. Some of the examples shown would only be valid for certain reports or files. The examples shown only show the Value column of the filtering panel, and always assume an Element name of "SAS" and an Operator of "EQ";

```
Select data from all SYSIDs with an 'X' in the third position: IF SUBSTR(SYSID,3,1) EQ 'X';
```

```
Ignore data for all volume serial numbers with a JES prefix: IF VOLSER EQ 'JES' THEN DELETE;
```

```
Select only data sets with a payroll or sales prefix: IF DSNAME EQ 'PAY' OR DSNAME EQ 'SALES';
```

```
Select only 3390 volumes for this analysis: IF DEVTYP EQ '3390';
```

```
Select only data sets occupying more than 100 tracks:
```

```
IF (FILE EQ 'DAA' OR FILE EQ '_VS') AND DAASPACA LT  
100 THEN DELETE;
```

If you have a complex series of SAS statements you want to include, but that cannot fit on the filtering screen, follow this procedure:

- o Create the SAS statements you want to include and place them in a data set, such as your Work Data Set.
- o Go to the MICF Options panel and select the option that allows "Additional Data Set Allocations." Use this panel to specify the name of the external data set containing the SAS statements, plus a ddname (any name you choose) that should be used to allocate it. When you have completed this option screen, the data set you specified will always be allocated when MICF executes an inquiry for you, and will always be allocated to the ddname you specified.
- o On the data filtering screen, use a SAS INCLUDE statement to include your SAS code. For example, assume you placed some SAS code in member MYSAS of your Work Data Set. Then you modified your MICF Options to always allocate your Work Data Set with a ddname of ADDSAS. You would then include this statement on your data filtering panel when you generate the desired StorageMate report:

```
%INCLUDE ADDSAS(MYSAS);
```

You could have multiple members containing SAS code defined for different reports, and then include the correct code when running each report.

3.10 How to Enable Neugents technology in CA MICS

The CA MICS Performance Manager, Capacity Planner, and StorageMate products include features that use CA's U.S. patented Neugents(R) technology.

Before you use Neugents, be sure your environment is ready. See these subsequent sections for more information:

- 1 - Requirements
- 2 - Activate the CA MICS Root Directory
- 3 - Enable Neugents Technology

3.10.1 Requirements

CA MICS delivers mainframe-based Neugents technology that requires C/C++ runtime features available in the language environment delivered with z/OS 1.3 and above. If you plan to use CA MICS product features that rely on Neugents technology, then CA MICS must be installed on a z/OS system at level 1.3 or above.

3.10.2 Activate the CA MICS Root Directory

To activate the CA MICS root directory, review Section 5.7 in the PIOM Guide and complete the activation checklist detailed in Section 5.7.3.

NOTE: If the CA MICS root directory has been enabled during the installation of another CA MICS product, you do not have to do it again.

3.10.3 Enable Neugents technology

Once you have activated the CA MICS root directory on USS, you can enable Neugents technology. By following the steps in this procedure, you will enable Neugents technology by delivering the Neugents technology DLLs to the correct directory under the CA MICS root directory.

To enable Neugents technology, complete the following steps:

- 1.____ Edit `sharedprefix.MICS.PARMS(JCLGENUC)` so that it contains the following:

```
MQRHFSU
```

Submit the job in `sharedprefix.MICS.CNTL(JCLGENUC)`. Ensure that there are no error messages in `MICSLOG` or `SYSTSPRT`, that the `MICSLOG` contains the normal termination message, `BAS10999I`, and that the job completes with a condition code of zero.

- 2.____ Unload Neugents technology HFS files to the CA MICS Root Directory.

Submit the following job:

```
'sharedprefix.MICS.CNTL(MQRHFSU)'
```

Ensure that there are no error messages in `MICSLOG` or `SYSTSPRT`, and that the job completes with a condition code of zero.

Chapter 4: REPORTS

This chapter describes the reports available through the reporting facilities of the StorageMate product. Separate sections are provided that describe applications of the reports from both SMS and non-SMS environments.

Figure 4-1 displays each inquiry available in CA MICS StorageMate arranged by storage administration category. The figure is intended to provide an easy reference guide to matching CA MICS StorageMate inquiries with specific storage management tasks.

Storage Administration Category	Report Name	Report ID
Problem Management	Unmanageable Data Sets	STGELA
	Volume Problems	STGECA
	Allocation Problems	STGECB
	New Data Set Allocation	STGECG
	Non-VSAM Wasted Space	STGECC
	VSAM Wasted Space	STGECE
	DASD/HSM Data Movement	STGECE
Data Set Standards	Data Set Names	STGEFA
	Non-VSAM Attributes	STGEFB
	VSAM Attributes	STGEFC
	GDG Management and Retention	STGEFE
Volume Grouping	Volume Group Configuration	STGEHA
	Volume Group Activity by Workload	STGEHB
	Volume Group Allocations by User	STGEHC
Space Availability	Deletion/Migration/Rlse Thresholds	STGEJA
	Space Management Funct Simulation	STGEJB
	Daily DASD Resource Usage	STGEJE
	Daily Application Resource Usage	STGERG
	Daily HSM DASD Usage	STGEJC

SMS Migration and Tracking	ACS Routine Definition	STGELB
	ACS Test Case Creation	STGELF
	ACS Testing Results	STGELG
	SMS Class Usage	STGELE
	SMS Implementation	STGELC
Performance Management	Data Set Performance	STGEPA
	Volume Configuration	STGEPB
	I/O Activity by Workload	STGEPC
	Daily ASTEX Cache Statistics	STGEPE
	ASTEX / Top Resource Consumers	STGEPF
Capacity Management	DASD-to-Tape Candidates	STGERA
	Tape-to-DASD Candidates	STGERB
	Space Allocation by User	STGERC
	Monthly DASD Resource Usage	STGERH
	Monthly Application Resource Usage	STGERE
	DASD/HSM Usage Snapshot	STGERF

Figure 4-1. StorageMate Reports List

This section contains the following topics:

[4.1 SMS Applications](#) (see page 71)

[4.2 Non-SMS Applications](#) (see page 76)

[4.3 Report Descriptions](#) (see page 79)

4.1 SMS Applications

This section presents each StorageMate report title with a brief description of its SMS-related applications. This section is organized to reflect panel flow within CA MICS StorageMate. Detailed information on each report is found in Section 4.3 of this guide; reports are listed alphabetically.

STORAGE ADMINISTRATION

PROBLEM MANAGEMENT

- Unmanageable Data Sets

Helps you identify data sets that cannot be managed by SMS.

- Volume Problems

Calculates actual and estimated VVDS sizes.

- Allocation Problems

After SMS is implemented, this report can be used to make sure abends are not occurring for managed users, and to alter SMS parameters if necessary.

- New Data Set Allocation

Allocation patterns detected by this report are useful when designing storage groups. After SMS is implemented, this report can be used to monitor new allocations based on storage class, management class, or storage group.

- Non-VSAM Wasted Space

- VSAM Wasted Space

Identifies groups with the worst space usage values, and helps determine efficient allocation values.

- DASD/HSM Data Movement

Allows the establishment and tuning of management class parameters relating to the migration of unused data sets.

DATA SET STANDARDS

- Data Set Names

Identifies common data set name patterns to assist in classifying the data sets into logical groups.

- Non-VSAM Attributes

- VSAM Attributes

Identifies and defines similar data sets that can be managed as one data class, and reveals patterns that can be used in ACS routines to make class assignments. Also helps determine the default space values for SMS data classes.

- GDG Management and Retention

Determines the number of generations of a GDG that are kept online and helps determine when to "roll off" old generations.

VOLUME GROUPING

- Volume Group Configuration

Plans and simulates SMS storage groups prior to SMS implementation; monitors those groups after implementation.

- Volume Group Activity by Workload

Simulates storage groups prior to SMS implementation, ensuring that I/O activity against those groups is distributed as expected; monitors those groups after implementation.

- Volume Group Allocations by User

Simulates storage groups prior to SMS implementation, ensuring that volumes with similar allocation patterns are grouped together; and monitors those groups after implementation.

SPACE AVAILABILITY

- Deletion/Migration/Release Thresholds

Monitors data to ensure that the volumes specified in the management classes are working as expected.

- Space Management Function Simulation

Helps ensure the values you specify in your management class definition will achieve the desired result.

- Daily DASD Resource Usage

Permits short-term monitoring of various capacity and performance statistics for SMS storage groups.

- Daily Application Resource Usage

Permits short-term monitoring of various capacity and usage statistics for each SMS storage class or management class.

- Daily HSM DASD Usage

Presents several statistics that measure the effectiveness of various HSM space management functions against SMS storage groups.

SMS MIGRATION AND TRACKING

- ACS Routine Definition

Identifies the data set pattern masks that should be coded in SMS ACS routines for the purpose of assigning class and group names for new data sets.

- ACS Test Case Creation

Allows faster and easier creation of test data used to validate the operation of the installation's ACS routines.

- ACS Testing Results

Allows quick, easy, and accurate evaluation of the test results from your ACS testing.

- SMS Class Usage

During SMS implementation, ensures that your SMS ACS routines assign class names as intended.

- SMS Implementation

Monitors the DFSMS conversion efforts made by individual

departments, groups, volumes, or systems.

PERFORMANCE MANAGEMENT

- Data Set Performance

Monitors response times for related groups of data sets and assists in planning storage classes.

- Volume Configuration

Helps define SMS storage classes and storage groups.

- I/O Activity by Workload

Groups volumes with similar workloads, which avoids performance problems caused by conflicting users.

- Daily ASTEX Cache Statistics

Identifies cache usage patterns helpful when defining Storage Classes and Storage Groups.

- ASTEX / Top Resource Consumers

Identifies the primary users of I/O resources, allowing for the definition of well-balanced Storage Groups.

CAPACITY MANAGEMENT

- DASD-to-Tape Candidates

- Tape-to-DASD Candidates

Identifies data sets that are candidates to be moved to more efficient storage locations based on their size and usage patterns.

- Space Allocation by User

Helps you group volumes with similar allocation patterns.

- Monthly DASD Resource Usage

- Monthly Application Resource Usage

Provides long-term capacity and performance information for the critical resources managed by SMS.

- DASD/HSM Usage Snapshot

Allows the monitoring of SMS class usage across both primary and HSM storage platforms.

4.2 Non-SMS Applications

This section presents each StorageMate report title with a brief description of its non-SMS-related applications. They are organized to reflect panel flow within CA MICS StorageMate. Detailed information on each report is found in Section 4.3 of this guide.

STORAGE ADMINISTRATION

PROBLEM MANAGEMENT

- Unmanageable Data Sets

Identifies data sets that may not be supported by automated storage management products until modified or converted.

- Volume Problems

Identifies problems with DASD volumes that could cause poor performance or availability.

- Allocation Problems

Identifies conditions that cause batch job abends related to DASD space shortages. Also reports the amount of CPU time expended in recovering from job abends related to DASD space shortages.

- New Data Set Allocation

Monitors the regular allocation of new data sets, and helps identify potential capacity and usage problems before they occur.

- Non-VSAM Wasted Space

- VSAM Wasted Space

Identifies users or departments that are wasting DASD storage space due to poor allocation practices.

- DASD/HSM Data Movement

Monitors the movement of data between primary DASD and HSM, identifying potential problems or abuses.

DATA SET STANDARDS

- Data Set Names

Reports the most common patterns in your data set names.

- Non-VSAM Attributes

- VSAM Attributes

Identifies common patterns of data set attributes, and provides a data set name node report and a space allocation range report for the data sets within a common group.

- GDG Management And Retention

Monitors your users' allocation and use of Generation Data Groups (GDGs).

The report allows analysis of generation data set allocation and usage patterns, and detects potential problems, such as duplicate or missing generations.

VOLUME GROUPING

- Volume Group Configuration

Summarizes performance and capacity information at the volume group level.

- Volume Group Activity by Workload

Monitors the I/O activity of each DASD volume group based on the EXCP counts stored by CA MICS in the BATWDA and HARDVA files.

- Volume Group Allocations by User

Summarizes how each DASD volume group is allocated by using account fields, data set name nodes, or a user-defined table.

SPACE AVAILABILITY

- Deletion/Migration/Release Thresholds

Identifies the age patterns of data sets by volume.

- Space Management Function Simulation

Represents the effect of space management actions on DASD volumes.

- Daily DASD Resource Usage

Permits the daily monitoring of storage resource usage at the volume, volume group, device type, or system level.

- Daily Application Resource Usage

Permits the daily monitoring of storage resource usage at the application level.

- Daily HSM DASD Usage

Displays information related to the effectiveness of HSM space management operations.

PERFORMANCE MANAGEMENT

- Data Set Performance

Estimates the average response time being received by each department or group.

- Volume Configuration

Summarizes performance and capacity information at the volume or device type level. Information may be reported based on the volume, device type, string or system.

- I/O Activity by Workload

Profiles the I/O activity of each DASD volume based on the workload EXCP counts stored by CA MICS in the BATWDA file.

- Daily ASTEX Cache Statistics

Allows the daily monitoring and tuning of DASD cache resources.

- ASTEX / Top Resource Consumers

Identifies the primary users of your critical storage resources, and isolates potential performance problems.

CAPACITY MANAGEMENT

- DASD-To-Tape Candidates

- Tape-To-DASD Candidates

Identifies data sets that are candidates to be moved to more efficient storage locations based on their size and usage patterns.

- Space Allocation by User

Summarizes DASD volume allocation by account field, data set node name, or by use of a user-defined table.

- Monthly DASD Resource Usage

- Monthly Application Resource Usage

Provides long-term capacity and performance information concerning your critical storage resources.

- DASD/HSM Usage Snapshot

Allows the monitoring of data distribution across both the primary and HSM storage platforms.

4.3 Report Descriptions

This section presents detailed descriptions of each report available in the StorageMate product. Reports are arranged alphabetically, and contain information on functions and features, as well as descriptions of the columns and one-page graphic samples of the report output. A dot in a report sample indicates that no value is present.

ACS Routine Definition / By Pattern (STGELB)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan
Tape library information (external)

Function: Allows users to define a series of data set name mask patterns, compare these patterns against their existing data set names, and report the results of these comparisons.

Provides a summary of each pattern, the number of matches that occurred for the pattern, and the space represented by the matching data sets.

Further information about the use of pattern files is available in the facility instructions at the end of this section.

Features: Many software products allow generic patterns or masks to be defined that are then used to allow operations on data sets whose names match the mask. This report helps users define, test, and tune those patterns so that software can be implemented faster and with fewer problems.

This analysis reads an external file of user-defined data set name patterns. This pattern file is then matched against each of the installation's data set names, with the results displayed on this report. The number of data sets that matched none of the patterns is also reported.

SMS

Issues: This report is useful for identifying the data set pattern masks that should be coded in SMS ACS routines for the purpose of assigning class and group names for new data sets. Once the pattern file is complete, a report option can be used to generate a starter or skeleton ACS routine.

COLUMN DESCRIPTIONS

Search Pattern

Identifies the search pattern that is summarized on this report line. Search patterns are defined

in a pattern file that you define and then read with this program. A value of '(NO MATCH)' in this column is used to represent data sets that did not match any of the existing patterns in the pattern file.

Construct Name Assigned

The name of the construct associated with this search pattern that was assigned to the data sets in this group. A construct name can, optionally, be associated with each search pattern in the pattern file. The concept of constructs allows many data sets with similar requirements to be grouped as one logical entity. For example, the following entries in the pattern file:

```
**..COBOL  SOURCE
**..ASM    SOURCE
**..PLI    SOURCE
```

would assign data set names ending in the suffixes COBOL, ASM and PLI to a common group known as 'SOURCE'. Because these data sets all probably represent source code, it is logical that they have similar requirements and can be treated as one entity. Depending on the type of ACS routine being written, the construct name could represent the name of a Data class, Storage class, Management class, or Storage Group. A value of '(NONE)' in this column indicates that no construct name was assigned, either because one was not defined in the pattern file, or because the data set matched none of the search patterns.

Data Sets Matching Pattern

The number of data sets within the input data that matched the search pattern represented by this summary line. A value of zero indicates a search pattern that did not match any data sets.

Percent Matching Pattern

The percentage of data sets within the input data that matched the search pattern represented by this summary line. A value of zero indicates a search pattern that did not match any data sets. Note that data sets excluded from the analysis by filtering criteria are not included when calculating this value.

Megabytes Allocated

The amount of storage space, expressed in millions of bytes (megabytes), allocated by the data sets matching this search pattern.

Percent Space Allocated

The percentage of the total storage space allocated by the data sets matching this search pattern, expressed in millions of bytes (megabytes). This is calculated by dividing the space allocated for this group by the total space allocated by all data sets in the analysis. Note that data sets excluded from the analysis by filtering criteria are not included when calculating this value.

Megabytes Used

The amount of storage space, expressed in millions of bytes (megabytes), used by the data sets matching this search pattern. Data sets included in the analysis from tape or HSM will always have this value equal to the allocated value (unless you specify the PRIMARY option for HSM space reporting).

Percent Space Used

The percentage of the total storage space used by the data sets matching this search pattern, expressed in millions of bytes (megabytes). This is calculated by dividing the space used for this group by the total space used by all data sets in the analysis. Note that data sets excluded from the analysis by filtering criteria are not included when calculating this value.

ACS ROUTINE DEFINITION (STGELB)							
CA MICS StorageMate - PATTERN REPORT							
ABC CORPORATION							
----- System Identifier=SYSA -----							
Search Pattern	Construct Name Assigned	Data Sets Matching Pattern	Percent Matching Pattern	Megabytes Allocated	Percent Space Allocated	Megabytes Used	Percent Space Used
TMPVSM,**	DCVSAMT	34	1	146.4	2	110.7	2
**.*LIST	DCLIST	298	9	179.9	3	151.4	3
**.*COBOL	DCSOURCE	0	0	0.0	0	0.0	0
**.*ASM	DCSOURCE	58	2	194.7	3	170.5	4
**.*PLI	DCSOURCE	0	0	0.0	0	0.0	0
**.*CNTL	DCSOURCE	263	8	196.9	3	158.7	3
**.*TEXT	DCDOCS	251	7	475.2	8	424.9	9
**.*ISP%LIB	DCISPF	183	5	607.9	10	540.4	12
(NO MATCH)	(NONE)	2296	68	4136.2	70	2984.6	66
-----	-----	-----	-----	-----	-----	-----	-----
SYSID		3383		5937.2		4541.3	
		=====		=====		=====	
		3383		5937.2		4541.3	

Figure 4-2. ACS Routine Definition / by Pattern

ACS Routine Definition / By Group (STGELB)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan
Tape library information (external)

Function: This report allows users to define a series of data set name mask patterns, compare those patterns against their existing data set names, assign a group name when a match occurs, and then summarize the data based on that group name. This particular report gives a summary of each of the assigned groups, including the number of data sets that would exist in the group, and the amount of space allocated and used by those data sets.

Further information about the use of pattern files is available in the facility instructions at the end of this section.

Benefit: Many software products allow generic patterns or masks to be defined that are then used to assign processing groups to the data set based on the results of the pattern matching. This report will help users define, test, and tune those patterns and groups so that software can be implemented faster and with fewer problems. These types of patterns and groups are very common in both automated data management software and in data security products.

Features: Because of the complexity of pattern testing, an assembler language format routine is called to do the actual testing. This routine allows single character masks (%), multiple character masks (*) and multiple node masks (**).

SMS

Issues: This report is useful for identifying the data set pattern masks that should be coded in SMS ACS routines for the purpose of assigning class and group names for new data sets. Once the pattern file is complete, a report option can be used to generate a starter or skeleton ACS routine.

COLUMN DESCRIPTIONS

Construct Name Assigned

The name of the construct for which this summarization applies. A construct name may, optionally, be associated with each search pattern in the pattern file. The concept of constructs allows many data sets with similar requirements to be grouped as one logical entity. For example, the following entries in the pattern file:

```
** .COBOL  SOURCE
** .ASM    SOURCE
** .PLI    SOURCE
```

would assign data set names ending in the suffixes COBOL, ASM and PLI to a common group known as 'SOURCE'. Because these data sets all probably represent source code, it is logical that they have similar requirements and can be treated as one entity. Depending on the type of ACS routine being written, the construct name could represent the name of a Data class, Storage class, Management class, or Storage Group. A value of '(NONE)' in this column indicates that no construct name was assigned, either because one was not defined in the pattern file, or because the data set matched none of the search patterns.

Data Sets Assigned to Group

The number of data sets within the input data that matched search pattern and were assigned the construct name represented by this summary line.

Percent Assigned to Group

The percentage of data sets within the input data that were assigned to the construct name represented by this summary line. Note that data sets excluded from the analysis by filtering criteria are not included when calculating this value.

Megabytes Allocated The amount of storage space, expressed in millions of bytes (megabytes), allocated by the data sets assigned to this construct.

Percent Space Allocated

The percentage of the total storage space allocated by the data sets assigned this construct, expressed in millions of bytes (megabytes). This is calculated by dividing the

space allocated for this group by the total space allocated by all data sets in the analysis. Note that data sets excluded from the analysis by filtering criteria are not included when calculating this value.

Megabytes Used

The amount of storage space, expressed in millions of bytes (megabytes), used by the data sets assigned to this construct. Note that data sets included from tape and HSM will always have this value equal to the allocated value (unless you use the PRIMARY option for HSM space reporting).

Percent Space Used

The percentage of the total storage space used by the data sets assigned to this construct, expressed in millions of bytes (megabytes). This is calculated by dividing the space used for this group by the total space used by all data sets in the analysis. Note that data sets excluded from the analysis by filtering criteria are not included when calculating this value.

ACS ROUTINE DEFINITION (STGELB) CA MICS StorageMate - GROUP REPORT ABC CORPORATION						
----- System Identifier=SYSA -----						
Construct Name Assigned	Data Sets Assigned to Group	Percent Assigned to Group	Megabytes Allocated	Percent Space Allocated	Megabytes Used	Percent Space Used
(NONE)	2296	68	4136.2	70	2984.6	66
DCDOCS	251	7	475.2	8	424.9	9
DCISPF	183	5	607.9	10	540.4	12
DCLIST	298	9	179.9	3	151.4	3
DCSOURCE	321	9	391.6	7	329.2	7
DCVSAMT	34	1	146.4	2	110.7	2
-----	-----		-----		-----	
SYSID	3383		5937.2		4541.3	
	=====		=====		=====	
	3383		5937.2		4541.3	

Figure 4-3. ACS Routine Definition / by Group

ACS Routine Definition / Detail (STGELB)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMMIG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan
Tape library information (external)

Function: Allows users to compare a series of data set name mask patterns against their existing data set names, and to report the results of these comparisons.

Provides a detailed listing of the data sets that match particular patterns and can be used to identify data sets that did not match any pattern.

Before using this facility, you must first create the pattern file that will be used for the comparison.

Further information about the use of pattern files is available in the facility instructions at the end of this section.

Features: Many software products allow generic patterns or masks to be defined that are then used to allow operations on data sets whose names match the mask. This report helps users define, test, and tune those patterns so that software can be implemented faster and with fewer problems.

This analysis reads an external file of data set patterns that have been defined by the user. This pattern file also indicates whether a detailed listing should be produced for all the data sets that match a particular pattern, and if detailed reporting is desired for those data sets that match none of the patterns.

SMS

Issues: This report is useful for identifying the data set pattern masks that should be coded in SMS ACS routines for the purpose of assigning class and group names for new data sets. Once the pattern file is complete, a report option can be used to generate a starter or skeleton ACS routine.

COLUMN DESCRIPTIONS

Search Pattern

Identifies the search pattern that matched the data set names that appear on this page of the report. A value of '(NO MATCH)' means that the data sets that appear in this section did not match any of the search patterns in the pattern file. Control over which data sets appear on the detailed report is set as part of the pattern file.

Data Set Name

The name of the data element.

Device Type

This value indicates the type of I/O device. Other possible values for this column include 'HSM-MIG' (migrated data sets), 'HSM-BACK' (HSM backup data sets), and 'TAPE'.

Volume Serial Number

The volume serial number of the storage medium. Values 'HSMVL1' (HSM migration level 1), 'HSMVL2' (HSM migration level 2), and 'HSMBAC' (HSM backup) are also valid.

Data Set Type

The organization of the data within the data set.

Date Created

The date that the data set was first allocated to a DASD volume.
This value will always be missing (.) for data sets on HSM Backup storage.

Date Last Used

The date that the data set was last referenced.
This value will always be missing (.) for data sets on HSM Backup storage.

Extents

The number of non-contiguous segments of space allocated to the data set on one volume. This value will be blank for tape or HSM-managed data sets. This value will be blank for data sets on tape and HSM storage.

Kilobytes Allocated

The amount of space allocated by this data set,

expressed in thousands of bytes (kilobytes). This is calculated by multiplying the tracks allocated by the maximum track capacity.

Kilobytes Used

The amount of space used by this data set, expressed in thousands of bytes (kilobytes). This is calculated by multiplying the tracks used by the maximum track capacity. For data sets on tape and HSM storage, this will always be the same value as Kilobytes Allocated (unless the PRIMARY reporting option is specified for HSM space reporting).

ACS ROUTINE DEFINITION (STGELB)								
CA MICS StorageMate - DETAIL								
ABC CORPORATION								
----- System Identifier=SYSA Search Pattern=**.ASM -----								
Data Set Name	Device Type	Volume Serial Number	Data Set Type	Date Created	Date Last Used	Extents	Kilobytes Allocated	Kilobytes Used
ATTACH.COMPOSIT.PDE3254.PSP.ASM	3380	LNR901	P0	11JUL90	14AUG07	1	95	95
ATTACH.DMMG4230.ASM	3380	LNR902	P0	24APR90	06JUL07	7	332	332
CHECK.BSGG.ASM	3380	LNR909	P0	01NOV89	16AUG07	2	2848	2374
CPSD.RE8701C.ASM	3380	LNR908	P0	17AUG89	17AUG07	1	5032	5032
CPSD.PSP8710C.ASM	3380	LNR900	P0	16AUG89	17AUG07	5	3750	3750
CPSD.PSP9004C.ASM	3380	LNR903	P0	19APR90	21AUG07	4	9020	9020
DAT003.EAS.ASM	3380	LNR906	P0	15NOV89	10AUG07	16	10777	3466
DAT009.MADMI.ASM	3380	LNR901	P0	10JAN90	08AUG07	3	760	760
DAT009.MASTER.DMI.ASM	3380	LNR909	P0	03OCT88	08AUG07	8	3561	3561
DAT009.MASTER.MADMI.ASM	3380	LNR909	P0	17FEB90	08AUG07	1	760	760
DAT068.PRIVATE.ASM	3380	LNR901	P0	15MAR90	13JUL07	1	47	47
DAT079.TEMP.ASM	3380	LNR901	P0	25JAN90	14MAY07	1	475	475
DAT101.MMB.ASM	3380	LNR901	P0	18DEC89	13JUL07	16	3940	285
DAT117.ASM	3380	LNR903	P0	16NOV89	07AUG07	2	1092	1092
DAT134.FLAG410.ASM	3380	LNR903	P0	12FEB90	22AUG07	3	2848	1757
DAT145.BBO.ASM	3380	LNR903	P0	08JUL90	08JUL07	2	902	902
DAT145.SSEG080.PSP.ASM	3380	LNR901	P0	02MAY90	21AUG07	16	760	142
DAT153.SEN3601.PSP.ASM	3380	LNR907	P0	25JUL89	16JUN07	1	47	47
DAT178.STDCMP.ASM	3380	LNR901	P0	28FEB90	10MAY07	2	855	855
DAT178.GTG4738.ASM	3380	LNR903	P0	08JAN90	14MAY07	1	427	427
DAT178.VTG4738.PSP.ASM	3380	LNR903	P0	05JAN90	14MAY07	1	47	47
DAT191.SRC.ASM	3380	LNR903	P0	01MAY90	27JUL07	1	760	760
DAT229.SAB4853.PSP.ASM	3380	LNR901	P0	02NOV89	10JUN07	1	332	332
FLAG.FLAGDONE.ASM	3380	LNR906	P0	25AUG89	19AUG07	4	29906	29242
FLAG.FLAGSTAGE.ASM	3380	LNR900	P0	24MAY90	14AUG07	1	1424	1282
FLAG.FLAG470.ASM	3380	LNR905	P0	22DEC89	17AUG07	6	33469	33422
FLGSHOW.PSP9004C.ASM	3380	LNR908	P0	01JUN90	17AUG07	1	47	47
GNNST.ARCTOQCA.HHK4950.ASM	3380	LNR901	P0	08JAN90	14MAY07	2	427	427
GNNST.ARCTOQCA.TEST.ASM	3380	LNR908	P0	08AUG90	.	1	47	47
GNNST.S6C.ASM	3380	LNR902	P0	14MAR90	16AUG07	6	8973	8973
GNNST.S6C.TEST.ASM	3380	LNR908	P0	29MAY90	29MAY07	1	47	47
PRODUCT.COMPOSIT.ABC0000.BETA.ASM	3380	LNR904	P0	16APR90	22AUG07	8	380	47
PRODUCT.PR09004U.TEST.ASM	3380	LNR903	P0	02MAY90	13JUN07	1	95	95
PRODUCT.PSP9004C.PSP.ASM	3380	LNR906	P0	19APR90	21AUG07	1	8356	95
DATA.COMPOSIT.SHC4299.PSP.ASM	3380	LNR907	P0	15JUN90	09JUL07	11	522	522
DATA.YAE6665.PSP.ASM	3380	LNR903	P0	20JUL90	21AUG07	1	712	190
DATA.MASTER.SP.ASM	3380	LNR901	P0	20JUL90	21AUG07	15	19891	19414
TESTING.DEV.PSP.ASM	3380	LNR900	P0	11JUN90	21AUG07	9	1709	1709
TESTING.JAE0270.ASM	3380	LNR901	P0	27FEB90	14MAY07	1	95	95
TESTING.HAE4276.ASM	3380	LNR908	P0	13AUG90	20AUG07	2	142	142

Figure 4-4. ACS Routine Definition / Detail

ACS Test Case Creation (STGELF)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMMIG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan
Tape library information (external)

Function: This facility obtains data about your existing data sets from up to five different input sources, and uses these data to build test cases for your ACS routines. These test cases can then be used as input to the ISMF ACS Testing facility to make sure your current or proposed ACS routines work correctly under a wide variety of conditions.

The output of this facility includes a summary report and a jobstream in your Work Data Set that can then be used to allocate a test case data set, and populate it with test cases constructed from your own data.

The ISMF ACS Testing facility can then process the test case library you have built, passing each test case to your ACS routines (either the active ones or test versions). The results of these tests will produce a report file that can then be analyzed by another CA MICS StorageMate facility--ACS Testing Results.

Further information about the use of the ACS Test Facility is available in the facility instructions at the end of this section.

Features: The ISMF ACS Testing facility allows you to build your own test cases, but the information has to be entered manually through ISPF, with each completed panel resulting in only one test case. This facility provides a mechanism that is faster and less prone to error.

Much of the information needed to describe each data set can be extracted directly from the CA MICS database. If your ACS routines are dependent upon information not in the database (such as the creating jobname of the data set), a secondary screen can be used to supply these values to all created test cases. Your own SAS code can also be inserted at a special exit point that receives

control after each test case has been constructed.

Because the number of data sets represented by even one cycle of CA MICS data may be prohibitive, two options have been included that reduce the number of test cases created. The Random option will select final test cases at a user-specified interval, such as every 5th case or every 20th case. The Subset option will include all data sets having unique characteristics, while excluding those similar to cases already selected. This option considers such factors as the data set name, data set size, accounting fields, and data set organization.

SMS

Issues: An important phase of any SMS conversion project is making sure your ACS routines are operating correctly and achieving their desired results. This problem is no less critical after SMS conversion, as regular maintenance to the ACS routines will still be necessary. Because of their critical nature, accuracy of ACS routines is just as important as the accuracy of any system exit. Testing your routines with a large number of test cases will help insure your success in this critical task.

COLUMN DESCRIPTIONS

SUMMARY OF TEST CASES CREATED IN MEMBER name OF
DATA SET dsname

This informational line identifies the Work Data Set (dsname) and member (name) containing the jobstream that will create and populate the test case library. This jobstream should be edited, if necessary, and then submitted to create the test case library.

CASES FROM DASD DATA (NON-VSAM)

A numeric value that represents the total number of test cases created from the CA MICS VCADAA file. VCADAA contains information about non-VSAM data sets that reside on primary (Level 0) DASD.

CASES FROM DASD DATA (VSAM)

A numeric value that represents the total number of test cases created from the CA MICS VCA_VS file. VCA_VS contains information about VSAM data sets that reside on primary (Level 0) DASD.

CASES FROM TAPE DATA

A numeric value that represents the total number of test cases created from data representing your tape library. This information is obtained from an external data set you provided. The MICSLOG report should identify the name of the data set that was used.

CASES FROM HSM MIGRATION DATA

A numeric value that represents the total number of test cases created from the CA MICS HSM MIG file. This file contains information about data sets that previously resided on DASD, but have migrated to either Level 1 or Level 2 HSM storage.

CASES FROM HSM BACKUP DATA

A numeric value that represents the total number of test cases created from the CA MICS HSBAC file. This file contains information about data sets that have one or more versions available on HSM backup storage.

TOTAL TEST CASES GENERATED

A numeric value that represents the total number of test cases created in the test case library. This number should be the sum of all the values described above.

ACS TEST CASE CREATION (STGELF)	
CA MICS StorageMate	
ABC CORPORATION	
SUMMARY OF TEST CASES CREATED IN MEMBER ACS0412 OF DATA SET STORADMN.MICSSTOR.CNTL	
CASES FROM DASD DATA (NON-VSAM)	742
CASES FROM DASD DATA (VSAM)	104
CASES FROM TAPE DATA	503
CASES FROM HSM MIGRATION DATA	379
CASES FROM HSM BACKUP DATA	5

TOTAL TEST CASES GENERATED	1733

Figure 4-5. ACS Test Case Creation

ACS Testing Results / Statistics (STGELG)

NOTE: Please review the "Procedures for using the ACS Test Facility" section before attempting to use this facility. This section is contained in the description of the ACS Test Case Creation facility.

Source: ACS Test Case Library
ACS Testing Output Report File

Function: This facility analyzes the results of the ACS routine testing process to make sure your ACS routines are working correctly. Another CA MICS StorageMate facility, the ACS Test Case Creation facility, is used to build a library of ACS test cases from your CA MICS data. That library is read by the ISMF ACS Test Facility, and each test case is passed to your production/test ACS Routine(s). The test results are written by ISMF to a report output file. This facility reads both the test case library and the testing output file, merges them, and provides a number of reports useful in evaluating the operation of the ACS routines.

This particular report provides a summary page of general statistics, useful for obtaining a quick impression of the overall testing results. Other reports may then be used to provide more detail or to explore anomalies.

Further information about the use of the ACS Test Facility is available in the facility instructions at the end of this section.

Features: Options allow you to restrict reporting to only those ACS routines specified. For example, when executing the ISMF Test Facility you can specify that all four types of ACS routines be executed (Storage Class, Management Class, Data Class, Storage Group), but when running reports you can only be interested in results of the Data Class routine.

Up to eight different reports may be generated by this facility, each showing a different view of the test results, or showing a different summarization. The reports can be selected in any

combination.

SMS

Issues: An important phase of any SMS conversion project is making sure your ACS routines are operating correctly and achieving their desired results. This problem is no less critical after SMS conversion, because regular maintenance to the ACS routines will still be necessary. Because of their critical nature, accuracy of ACS routines is just as important as the accuracy of any system exit. Using this facility to carefully analyze your test results will help insure your success in this critical task.

COLUMN DESCRIPTIONS

TOTAL NUMBER OF TEST CASES:

A numeric value indicating the total number of test cases that were processed by the ACS Test Facility.

STORAGE CLASS COUNTS:

The statistics following this heading apply to the testing of the storage class ACS routine. Refer to the 'Common Values' section below for a description of each of these items.

MANAGEMENT CLASS COUNTS:

The statistics following this heading apply to the testing of the management class ACS routine. Refer to the 'Common Values' section below for a description of each of these items.

DATA CLASS COUNTS:

The statistics following this heading apply to the testing of the data class ACS routine. Refer to the 'Common Values' section below for a description of each of these items.

STORAGE GROUP COUNTS:

The statistics following this heading apply to the testing of the storage group ACS routine. Refer to the 'Common Values' section below for a description of each of these items.

(Common Values)

ACS ROUTINE INVOCATION:

A numeric value indicating the total number of times this ACS routine was invoked in this test. The data class and storage class ACS routines should be called for every test case. The management class ACS routine should be called for every managed data set, with a few exceptions such as VIO data sets. The storage group ACS routine should be called for every managed data set.

RETURN CODE NOT ZERO:

A numeric value indicating the number of times the ACS routine completed with a non-zero return code. This condition would cause the allocation to fail if the routine were running in a production environment. The number in parentheses indicates the percentage of the total calls to this ACS routine that had this result.

VALID CLASS/GROUP VALUE ASSIGNED:

A numeric value indicating the number of calls to this ACS routine that resulted in a zero return code, and a valid (non-blank) class/group name being assigned. The number in parentheses indicates the percentage of the total calls to this ACS routine that had this result.

BLANK (NULL) VALUE ASSIGNED:

A numeric value indicating the number of calls to this ACS routine that resulted in a zero return code, but a null (blank) class/group name being assigned. The number in parentheses indicates the percentage of the total calls to this ACS routine that had this result.

ACS TESTING RESULTS (STGELG)	
CA MICS StorageMate - STATISTICS	
ABC CORPORATION	
TOTAL NUMBER OF TEST CASES:	416
STORAGE CLASS COUNTS:	
ACS ROUTINE INVOCATION:	416
RETURN CODE NOT ZERO:	5 (1%)
VALID CLASS VALUE ASSIGNED:	398 (96%)
BLANK (NULL) VALUE ASSIGNED:	13 (3%)
MANAGEMENT CLASS COUNTS:	
ACS ROUTINE INVOCATION:	396
RETURN CODE NOT ZERO:	0 (0%)
VALID CLASS VALUE ASSIGNED:	386 (97%)
BLANK (NULL) VALUE ASSIGNED:	10 (3%)
DATA CLASS COUNTS:	
ACS ROUTINE INVOCATION:	416
RETURN CODE NOT ZERO:	0 (0%)
VALID CLASS VALUE ASSIGNED:	358 (86%)
BLANK (NULL) VALUE ASSIGNED:	58 (14%)
STORAGE GROUP COUNTS:	
ACS ROUTINE INVOCATION:	398
RETURN CODE NOT ZERO:	0 (0%)
VALID GROUP VALUE ASSIGNED:	398 (100%)
BLANK (NULL) VALUE ASSIGNED:	0 (0%)

Figure 4-6. ACS Testing Results / Statistics

ACS Testing Results / Detail (STGELG)

Source: ACS Test Case Library
ACS Testing Output Report File

Function: This facility analyzes the results of the ACS routine testing process to make sure your ACS routines are working correctly. Another CA MICS StorageMate facility, the ACS Test Case Creation facility, is used to build a library of ACS test cases from your CA MICS data. That library is read by the ISMF ACS Test Facility, and each test case is passed to your production/test ACS Routine(s). The test results are written by ISMF to a report output file. This facility reads both the test case library and the testing output file, merges them, and provides a number of reports useful in evaluating the operation of the ACS routines.

This particular report provides a detailed listing showing complete information about the results of each test case processed. It provides the most in-depth information available when researching a problem or verifying ACS routine operation.

Further information about the use of the ACS Test Facility is available in the facility instructions at the end of this section.

Features: Options allow you to restrict reporting to only those ACS routines specified. For example, when executing the ISMF Test Facility you may specify that all four types of ACS routines be executed (Storage Class, Management Class, Data Class, Storage Group), but when running reports you may only be interested in results of the Data Class routine.

Up to eight different reports may be generated by this facility, each showing a different view of the test results, or showing a different summarization. The reports can be selected in any combination.

SMS

Issues: An important phase of any SMS conversion project is making sure your ACS routines are operating correctly and achieving their desired results. This problem is no less critical after SMS conversion, as regular maintenance to the ACS

routines will still be necessary. Because of their critical nature, accuracy of ACS routines is just as important as the accuracy of any system exit. Using this facility to carefully analyze your test results will help ensure your success in this critical task.

COLUMN DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the system to which this data applies. The System Identifier (SYSID) is copied from the CA MICS data used in creating the test cases.

Maximum Return Code

A numeric value indicating the maximum return code received from any of the ACS routines called on behalf of this test case. This column will be blank if the return code was zero. Examination of the next four columns (class and group names) should reveal the error codes presented by individual routines.

Storage Class Name

The name of the storage class assigned to this data set by the storage class ACS routine. A BLANK value indicate that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a non-zero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Manage. Class Name

The name of the management class assigned to this data set by the management class ACS routine. A BLANK value indicate that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a non-zero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Storage Group Name

The name of the storage group assigned to this data set by the storage group ACS routine. A BLANK value indicate that the ACS routine was not invoked

for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a non-zero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Data Class Name

The name of the data class assigned to this data set by the data class ACS routine. A BLANK value indicate that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a non-zero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Data Set Name

The name of the data set passed to the testing routine for this test case.

Test Case Member

Identifies the test case library member (PDS member) that was passed to the ACS test routine. You may want to browse this member in the test case library to identify the actual values passed to the ACS routine. This is especially valuable when investigating errors or improper results in the ACS routines.

Size (KB)

The calculated size of the identified data set as passed to the ACS routine. The value is represented in kilobytes (thousands of bytes). This value is important because the ACS routine often makes decisions based on it.

Input Type

The location of the original data set upon which this test case was based. This value is supplied by the ACS Test Case Creation function, based on the original source of the data. Valid values are DASD (VCADAA file), VSAM (VCA_VS file), TAPE (External), MIGRATE (HSM MIG file), or BACKUP (HSBBAC file).

ACS TESTING RESULTS (STGELG)									
CA MICS StorageMate - DETAIL									
ABC CORPORATION									
----- SYSID=SYSA -----									
Maximum Return Code	Storage Class Name	Manage. Class Name	Storage Group Name	Data Class Name	Data Set Name	Test Case Member	Size (KB)	Input Type	
	STANDARD	STANDARD	TSOWRK	(NONE)	BACK01.DATASETS.SETS5.LIST	@BACA002	18.4	MIGRATE	
	STANDARD	LARGE	LARGE GP	DCPS	BKUP0F1.ARM.R.TEST.DATA	@BKUA001	100409.4	TAPE	
	STANDARD	STANDARD	TSOWRK	(NONE)	BSX.V22.PX.SOURCE	@BSXA002	510.0	MIGRATE	
	STANDARD	STANDARD	TSOWRK	DCPS	CAPPLN.COMPOSED.CAPTDTA.DATA	@CAPA003	226.7	DASD	
	STANDARD	STANDARD	TSOWRK	(NONE)	CAPPLN.DTESTLIB.ISPTLIB	@CAPA004	2.0	MIGRATE	
	(NONE)			DCPS	CSUT.BDBFR250.BDB.CLOSED.ADJUST.CNTL	@CSUA001	712.1	DASD	
	STANDARD	(NONE)	TSOWRK	DCPS	CSUT.PST8904A.CNTL	@CSUA003	793.3	DASD	
	(NONE)			DCPS	CSUT.PST8910A.INCL	@CSUA004	1139.4	DASD	
	STANDARD	STANDARD	TSOWRK	(NONE)	CSUT.PST8910A.TEST120.CNTL	@CSUA005	73.7	MIGRATE	
	STANDARD	STANDARD	TSOWRK	DCDA	CSUT.PST8910A.MONFILE	@CSUA006	4747.6	DASD	
	(NONE)			DCPS	DA1TEST.TSOWRK.PROCLIB	@DA1A001	56.7	DASD	
	(NONE)			DCPS	DA1TEST.MAST.USER.PSTT.LIST	@DA1A002	2.0	MIGRATE	
	STANDARD	STANDARD	TSOWRK	DCPS	DXBAAA.TST4455.MR	@DXBA047	47.5	DASD	
8	*RC=0008			DCPS	DXBBBB.LDFGEDS.PT.TEXTCNTL	@DXBA048	849.9	DASD	
8	*RC=0008			DCPS	DXBCCC.MAST.USER.BATCH1	@DXBA049	47.5	DASD	
	STANDARD	STANDARD	TSOWRK	(NONE)	DXBDDD.CICSTST.TEMPDATA	@DXBA068	20.5	MIGRATE	
12	*RC=0012			DCDA	DXBEEE.BXX302.DISK.SUM.DATA.Y002C001	@DXBA069	2848.4	DASD	
	STANDARD	STANDARD	TSOWRK	DCPS	DXBFFF.RRR.SMFDATA	@DXBA071	849.9	DASD	
	STANDARD	STANDARD	TSOWRK	DCPS	DXBGGG.TSTPROBS.SOURCE	@DXBA072	902.0	DASD	
12	*RC=0012			DCPS	DXBHHH.TST4000.PRT.PD.TEXT	@DXBA073	113.3	DASD	
8	*RC=0008			DCPS	DXBIII.PRODUCT.SMF	@DXBA074	712.1	DASD	
	STANDARD	NOMIG	TSOWRK	DCPS	SYS2.SMFBKUP.WEEKLY1	@SYSA012	134766.6	TAPE	
	STANDARD	NOMIG	TSOWRK	DCPS	SYS2.SMFBKUP.WEEKLY2	@SYSA013	135053.3	TAPE	
	VIODA		VIODA	DCPS	SYS3.XX.ISPPLIB	@SYSA014	4249.8	DASD	
	VIODA		VIODA	DCVS	SYS3.ZZTEST1.CXX.DATA	@SYSA015	39664.1	VSAM	
-----	SYSID						428816.2	-----	
							=====		
N=	25								

Figure 4-7. ACS Testing Results / Detail

ACS Testing Results / Exceptions (STGELG)

Source: ACS Test Case Library
ACS Testing Output Report File

Function: This facility analyzes the results of the ACS routine testing process to make sure your ACS routines are working correctly. Another CA MICS StorageMate facility, the ACS Test Case Creation facility, is used to build a library of ACS test cases from your CA MICS data. That library is read by the ISMF ACS Test Facility, and each test case is passed to your production/test ACS Routine(s). The test results are written by ISMF to a report output file. This facility read both the test case library and the testing output file, merges them, and provides a number of reports useful in evaluating the operation of the ACS routines.

The exception reports identify conditions that may indicate an error in the operation of the ACS routine. Two different exception reports may be requested. One report lists test cases that received a non-zero return code from the ACS routine, while the other report indicates test cases that received a blank (null) value back from one or more of the ACS routines.

Further information about the use of the ACS Test Facility is available in the facility instructions at the end of this section.

Features: Options allow you to restrict reporting to only those ACS routines specified. For example, when executing the ISMF Test Facility you can specify that all four types of ACS routines be executed (Storage Class, Management Class, Data Class, Storage Group), but when running reports you can only be interested in results of the Data Class routine.

Up to eight different reports can be generated by this facility, each showing a different view of the test results, or showing a different summarization. The reports can be selected in any combination.

SMS

Issues: An important phase of any SMS conversion project is making sure your ACS routines are operating

correctly and achieving their desired results. This problem is no less critical after SMS conversion, as regular maintenance to the ACS routines will still be necessary. Because of their critical nature, accuracy of ACS routines is just as important as the accuracy of any system exit. Using this facility to carefully analyze your test results will help ensure your success in this critical task.

COLUMN DESCRIPTIONS

TITLE 4

The fourth title line of the report indicates which exception report is being produced. This title will be either 'TEST CASES WITH NON-ZERO RETURN CODES' or 'TEST CASES WITH NULL VALUES ASSIGNED'. Test cases will appear on these reports if ANY of the ACS routines called for a test case returns a non-zero value or a null assignment.

SYSID

This value appears at the top of each page and identifies the system to which this data applies. The System Identifier (SYSID) is copied from the CA MICS data used in creating the test cases.

Maximum Return Code

A numeric value indicating the maximum return code received from any of the ACS routines called on behalf of this test case. This column will be blank if the return code was zero. Examination of the next four columns (class and group names) should reveal the error codes presented by individual routines.

Storage Class Name

The name of the storage class assigned to this data set by the storage class ACS routine. A BLANK value indicates that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a nonzero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Manage. Class Name

The name of the management class assigned to this

data set by the management class ACS routine. A BLANK value indicate that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a nonzero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Storage Group Name

The name of the storage group assigned to this data set by the storage group ACS routine. A BLANK value indicates that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a nonzero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Data Class Name

The name of the data class assigned to this data set by the data class ACS routine. A BLANK value indicates that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a nonzero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Data Set Name

The name of the data set passed to the testing routine for this test case.

Test Case Member

Identifies the test case library member (PDS member) that was passed to the ACS test routine. You may want to browse this member in the test case library to identify the actual values passed to the ACS routine. This is especially valuable when investigating errors or improper results in the ACS routines.

Size (KB)

The calculated size of the identified data set as passed to the ACS routine. The value is represented in kilobytes (thousands of bytes). This value is important as the ACS routine often makes decisions based on it.

Input Type

The location of the original data set upon which this test case was based. This value is supplied by the ACS Test Case Creation function, based on the original source of the data. Valid values are DASD (VCADAA file), VSAM (VCA_VS file), TAPE (External), MIGRATE (HSMIMG file), or BACKUP (HSBBAC file).

ACS TESTING RESULTS (STGELG) CA MICS StorageMate - EXCEPTIONS ABC CORPORATION TEST CASES WITH NULL VALUES ASSIGNED									
----- SYSID=SYSA -----									
Maximum Return Code	Storage Class Name	Manage. Class Name	Storage Group Name	Data Class Name	Data Set Name	Test Case Member	Size (KB)	Input Type	
	STANDARD	STANDARD	TSOWRK	(NONE)	ATA.TEST2.LOAD	@ATAA002	364.5	MIGRATE	
	STANDARD	STANDARD	TSOWRK	(NONE)	BACKIT.DATASETS.SETS5.LIST	@BACA002	18.4	MIGRATE	
	STANDARD	STANDARD	TSOWRK	(NONE)	BSS.V22.P1.SOURCE	@BSSA002	510.0	MIGRATE	
	STANDARD	STANDARD	TSOWRK	(NONE)	BTEST.PROD.AAAFMCF	@BTEA002	739.3	MIGRATE	
	STANDARD	STANDARD	TSOWRK	(NONE)	CAPPLN.C.PXX.PROCLIB	@CAPA001	65.5	MIGRATE	
	STANDARD	STANDARD	TSOWRK	(NONE)	CAPPLN.DEVT1111.ISPTLIB	@CAPA004	2.0	MIGRATE	
	(NONE)			DCPS	CST.RCP.COMPLOAD.PREPROC	@CSTA001	94.9	DASD	
	STANDARD	STANDARD	TSOWRK	(NONE)	CSTPLN.ACTIVET.PST.DOC.TEXT	@CSTA002	149.5	MIGRATE	
	(NONE)			DCPS	CSUT.BXXXXR250.XXX.CLOSED.ADJUST.CNTL	@CSUA001	712.1	DASD	
	STANDARD	(NONE)	TSOWRK	(NONE)	CSUT.BXXXXR302.DOCFIX	@CSUA002	10.2	MIGRATE	
	STANDARD	(NONE)	TSOWRK	DCPO	CSUT.PXXXX04U.CNTL	@CSUA003	793.3	DASD	
	(NONE)			DCPO	CSUT.PXXXX10C.INCLLIB	@CSUA004	1139.4	DASD	
	STANDARD	STANDARD	TSOWRK	(NONE)	CSUT.PXXXX10C.PSTTEST.CNTL	@CSUA005	73.7	MIGRATE	
	(NONE)			DCPS	WDDA.PROD.TCORP.XXXBEXT.BKUP	@WDDA004	9113.3	TAPE	
	(NONE)			DCPS	WDDA.WDDADBTC.STEP87.PRCOPY	@WDDA005	32.8	TAPE	
-----							-----		
SYSID							13818.9		
							=====		
N=	15								

Figure 4-8. ACS Testing Results / Exceptions

ACS Testing Results / Analysis (STGELG)

Source: ACS Test Case Library
ACS Testing Output Report File

Function: This facility analyzes the results of the ACS routine testing process to make sure your ACS routines are working correctly. Another CA MICS StorageMate facility, the ACS Test Case Creation facility, is used to build a library of ACS test cases from your CA MICS data. That library is read by the ISMF ACS Test Facility, and each test case is passed to your production/test ACS Routine(s). The test results are written by ISMF to a report output file. This facility reads both the test case library and the testing output file, merges them, and provides a number of reports useful in evaluating the operation of the ACS routines.

The storage class ACS routine has a dual function, because it must decide whether a data set should be managed as well as determining the class to be assigned to it. The analysis reports are useful in making sure the first function is working correctly. One analysis report will show all data sets that would be managed, while the other shows those that would not be managed. These reports can only be requested if you specified that the storage class ACS routine be invoked as part of the testing process.

Further information about the use of the ACS Test Facility is available in the facility instructions at the end of this section.

Features: Options allow you to restrict reporting to only those ACS routines specified. For example, when executing the ISMF Test Facility you can specify that all four types of ACS routines be executed (Storage Class, Management Class, Data Class, Storage Group), but when running reports you may only be interested in results of the Data Class routine.

Up to eight different reports may be generated by this facility, each showing a different view of the test results, or showing a different summarization. The reports can be selected in any combination.

SMS

Issues: An important phase of any SMS conversion project is making sure your ACS routines are operating correctly and achieving their desired results. This problem is no less critical after SMS conversion, as regular maintenance to the ACS routines will still be necessary. Because of their critical nature, accuracy of ACS routines is just as important as the accuracy of any system exit. Using this facility to carefully analyze your test results will help ensure your success in this critical task.

COLUMN DESCRIPTIONS

TITLE 4

The fourth title line of the report indicates which analysis report is being produced. This title will be either 'MANAGED DATA SETS' or 'UNMANAGED DATA SETS'. If a test case is assigned a blank storage class, it will not be managed. A non-blank value, as well as a return code of zero, indicates SMS will manage the data set.

SYSID

This value appears at the top of each page and identifies the system to which this data applies. The System Identifier (SYSID) is copied from the CA MICS data used in creating the test cases.

Maximum Return Code

A numeric value indicating the maximum return code received from any of the ACS routines called on behalf of this test case. This column will be blank if the return code was zero. Examination of the next four columns (class and group names) should reveal the error codes presented by individual routines.

Storage Class Name

The name of the storage class assigned to this data set by the storage class ACS routine. A BLANK value indicates that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. This null value should never be present for a managed data set, but should always be present for an unmanaged one. A value of '*RC=nnnn' indicates a non-zero return code from

the routine, where 'nnnn' is the numeric representation of the return code.

Manage. Class Name

The name of the management class assigned to this data set by the management class ACS routine. A BLANK value indicates that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a non-zero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Storage Group Name

The name of the storage group assigned to this data set by the storage group ACS routine. A BLANK value indicates that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a non-zero return code from the routine where 'nnnn' is the numeric representation of the return code.

Data Class Name

The name of the data class assigned to this data set by the data class ACS routine. A BLANK value indicates that the ACS routine was not invoked for this data set. A value of '(NONE)' indicates that the ACS routine was called, but no value was assigned. A value of '*RC=nnnn' indicates a non-zero return code from the routine, where 'nnnn' is the numeric representation of the return code.

Data Set Name

The name of the data set passed to the testing routine for this test case.

Test Case Member

Identifies the test case library member (PDS member) that was passed to the ACS test routine. You may want to browse this member in the test case library to identify the actual values passed to the ACS routine. This is especially valuable when investigating errors or improper results in the ACS routines.

Size (KB)

The calculated size of the identified data set as passed to the ACS routine. The value is represented in kilobytes (thousands of bytes). This value is important as the ACS routine often makes decisions based on it.

Input Type

The location of the original data set upon which this test case was based. This value is supplied by the ACS Test Case Creation function, based on the original source of the data. Valid values are DASD (VCADAA file), VSAM (VCA_VS file), TAPE (External), MIGRATE (HSM MIG file), or BACKUP (HSBBAC file).

ACS TESTING RESULTS (STGELG) CA MICS StorageMate - ANALYSIS ABC CORPORATION UNMANAGED DATA SETS									
----- SYSID=SYSA -----									
Maximum Return Code	Storage Class Name	Manage. Class Name	Storage Group Name	Data Class Name	Data Set Name	Test Case Member	Size (KB)	Input Type	
(NONE)				DCPS	CSA.RCP.COMPLOAD.PREPROC	@CSAA002	94.9	DASD	
(NONE)				DCPS	CSUT.AXXXX250.XXX.CLOSED.ADJUST.CNTL	@CSUA001	712.1	DASD	
(NONE)				DCPO	CSUT.PSTEST0C.INCLLIB	@CSUA004	1139.4	DASD	
(NONE)				DCPO	DA1TES1.TS0TST.PROCLIB	@DA1A001	56.7	DASD	
(NONE)				DCPS	DA1TES2.MAST.USER.PSTS.LIST	@DA1A002	2.0	MIGRATE	
(NONE)				DCPS	DA2TEST.NDXX.D01214.T161007	@DA2A002	2.0	MIGRATE	
(NONE)				DCPS	DBRT.DALYTRAN.BACKUP.G1594V00	@DBRA001	7858.4	TAPE	
(NONE)				DCPS	DBRT.PACK.DBXX01.G1688V00	@DBRA002	206257.0	TAPE	
(NONE)				DCPS	DDAT.DDATEST4.TRNORT.BKOKD.G3188V00	@DDAA001	21533.2	TAPE	
(NONE)				DCPO	DSITEST.ASM.CLASS	@DSIA006	192.5	MIGRATE	
(NONE)				(NONE)	DXBXXX.BXXX302.MODL5CB	@DXBA040	0.0	DASD	
(NONE)				DCPS	WDDA.PROD.TEST1.XXXBEXT.BKUP	@WDDA004	9113.3	TAPE	
(NONE)				DCPS	WDDA.TESTDBTC.STEP21.PRTEST.G0002V00	@WDDA005	32.8	TAPE	
-----							246994.4		
SYSID							=====		
							246994.4		
N= 13									

Figure 4-9. ACS Testing Results / Analysis

ACS Testing Results / Group (STGELG)

Source: ACS Test Case Library
ACS Testing Output Report File

Function: This facility analyzes the results of the ACS routine testing process to make sure your ACS routines are working correctly. Another CA MICS StorageMate facility, the ACS Test Case Creation facility, is used to build a library of ACS test cases from your CA MICS data. That library is read by the ISMF ACS Test Facility, and each test case is passed to your production/test ACS Routine(s). The test results are written by ISMF to a report output file. This facility reads both the test case library and the testing output file, merges them, and provides a number of reports useful in evaluating the operation of the ACS routines.

The group reports summarize the results of each ACS routine, showing the various class/group names that were assigned by the routine, and the number of data sets assigned to each of those names. Space values are also reported, to assist in the space planning needed for the various groups.

Further information about the use of the ACS Test Facility is available in the facility instructions at the end of this section.

Features: Options allow you to restrict reporting to only those ACS routines specified. For example, when executing the ISMF Test Facility you can specify that all four types of ACS routines be executed (Storage Class, Management Class, Data Class, Storage Group), but when running reports you can only be interested in results of the Data Class routine.

Up to eight different reports may be generated by this facility, each showing a different view of the test results, or showing a different summarization. The reports can be selected in any combination.

SMS

Issues: An important phase of any SMS conversion project is making sure your ACS routines are operating correctly and achieving their desired results. This problem is no less critical after SMS

conversion, as regular maintenance to the ACS routines will still be necessary. Because of their critical nature, accuracy of ACS routines is just as important as the accuracy of any system exit. Using this facility to carefully analyze your test results will help insure your success in this critical task.

COLUMN DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the system to which this data applies. The System Identifier (SYSID) is copied from the CA MICS data used in creating the test cases.

Storage Class Name

Management Class Name

Data Class Name

Storage Group Name

The first column of the report will have one of these titles, indicating the ACS routine to which the statistics on the report apply. A separate report will be produced for each ACS routine for which test data was found in the input report file. The values under the column heading indicate the various names assigned by the ACS routine. A value of '(NONE)' indicates that no class/group name was assigned by the ACS routine to the test cases in that group. A value of '*RC=nnnn' indicates a non-zero return code from the ACS routine, where 'nnnn' is the numeric representation of the return code.

Number of Data Sets

A numeric value indicating the number of test cases assigned by the ACS routine to this class or group name.

Percent of Data Sets

The percentage of all test cases for which this ACS routine was called that were assigned to this particular class or group.

Space Allocated M-Bytes

The amount of space, reported in megabytes, that is being allocated by the data sets that were assigned to this particular class or group. This value is calculated using data set allocation statistics extracted from CA MICS at the time the test cases

were created.

Percent of Megabytes

Based on the total megabytes allocated by all the data sets for which this ACS routine was invoked, the percentage of that space allocated by the data sets assigned to this particular class or group.

Minimum Size K-Bytes

The smallest amount of space allocated by any one of the data sets assigned to this class or group, expressed in kilobytes.

Maximum Size K-Bytes

The largest amount of space allocated by any one of the data sets assigned to this class or group, expressed in kilobytes.

Average Size K-Bytes

The average amount of space allocated by the data sets assigned to this class or group, expressed in kilobytes. This is calculated by dividing the total space allocated for all data sets in the group by the number of data sets assigned to the group.

ACS TESTING RESULTS (STGELG)								
CA MICS StorageMate - GROUP								
ABC CORPORATION								
----- SYSID=SYSA -----								
Storage Class Name	Number of Data Sets	Percent of Data Sets	Space Allocated M-Bytes	Percent of Megabytes	Minimum Size K-Bytes	Maximum Size K-Bytes	Average Size K-Bytes	
(NONE)	13	3.1	247.0	1.2	0	206257	19000	
STANDARD	396	95.2	20372.1	98.6	0	774971	51445	
VIOTEST	2	0.5	43.9	0.2	4250	39664	21957	

SYSID	411		20663.0					
----- SYSID=SYSA -----								
Management Class Name	Number of Data Sets	Percent of Data Sets	Space Allocated M-Bytes	Percent of Megabytes	Minimum Size K-Bytes	Maximum Size K-Bytes	Average Size K-Bytes	
(NONE)	10	2.5	4.4	0.0	6	1729	442	
LARGE	131	33.1	18572.5	91.2	2136	774971	141775	
NOMIGRT	12	3.0	1476.3	7.2	38756	135201	123021	
STANDARD	243	61.4	318.9	1.6	0	42496	1312	

SYSID	396		20372.1					
----- SYSID=SYSA -----								
Data Class Name	Number of Data Sets	Percent of Data Sets	Space Allocated M-Bytes	Percent of Megabytes	Minimum Size K-Bytes	Maximum Size K-Bytes	Average Size K-Bytes	
(NONE)	58	13.9	52.3	0.3	0	25393	902	
DCDA	14	3.4	128.8	0.6	95	33230	9200	
DCPO	92	22.1	101.6	0.5	2	28672	1104	
DCDA	236	56.7	20295.5	98.2	0	774971	85998	

SYSID	400		20588.2					
----- SYSID=SYSA -----								
Storage Group Name	Number of Data Sets	Percent of Data Sets	Space Allocated M-Bytes	Percent of Megabytes	Minimum Size K-Bytes	Maximum Size K-Bytes	Average Size K-Bytes	
CUSTSUP	5	1.3	2.9	0.0	57	1247	578	
LARGECP	131	32.9	18572.5	91.0	2136	774971	141775	
TSOWRK	260	65.3	1796.7	8.8	0	135201	6910	

SYSID	396		20373.1					

Figure 4-10. ACS Testing Results / Group

ACS Testing Results / Owner (STGELG)

Source: ACS Test Case Library
ACS Testing Output Report File

Function: This facility analyzes the results of the ACS routine testing process to make sure your ACS routines are working correctly. Another CA MICS StorageMate facility, the ACS Test Case Creation facility, is used to build a library of ACS test cases from your CA MICS data. That library is read by the ISMF ACS Test Facility, and each test case is passed to your production/test ACS Routine(s). The test results are written by ISMF to a report output file. This facility reads both the test case library and the testing output file, merges them, and provides a number of reports useful in evaluating the operation of the ACS routines.

The owner report produces a consolidated analysis of the operation of all four ACS routines, summarized by the application associated with each of the test case data sets. This provides an excellent method of cross-referencing the operation of all ACS routines, making sure they will operate correctly with each different application's data.

Further information about the use of the ACS Test Facility is available in the facility instructions at the end of this section.

Features: Options allow you to restrict reporting to only those ACS routines specified. For example, when executing the ISMF Test Facility you can specify that all four types of ACS routines be executed (Storage Class, Management Class, Data Class, Storage Group), but when running reports you can only be interested in results of the Data Class routine.

Up to eight different reports can be generated by this facility, each showing a different view of the test results, or showing a different summarization. The reports can be selected in any combination.

The application value for each data set is set when the test cases are built by the ACS Test Case Creation facility within CA MICS StorageMate. Options include application assignment based on one

or more nodes of the data set name, one or more of the CA MICS accounting fields, or using a user-defined table that associates data set prefixes with an application name.

SMS

Issues: An important phase of any SMS conversion project is making sure your ACS routines are operating correctly and achieving their desired results. This problem is no less critical after SMS conversion, as regular maintenance to the ACS routines will still be necessary. Because of their critical nature, accuracy of ACS routines is just as important as the accuracy of any system exit. Because most ACS routines will probably be written to assign critical classes to only selected applications, this report will be useful in ensuring those assignments are working correctly.

COLUMN DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the system to which this data applies. The System Identifier (SYSID) is copied from the CA MICS data used in creating the test cases.

Owner

Identifies the application or owner of the data sets in this group. This value is assigned to each data set when the test cases are created by the ACS Test Case Creation facility of CA MICS StorageMate. Refer to that facility for details about options for establishing data set ownership. All of the values in this report will apply to the data sets that were assigned to this particular owner.

Storage Class Name

The name of the class assigned to the data sets in this group by the storage class ACS routine. Note that a summary line will appear for each unique combination of storage class, management class, data class, and storage group. The report groups are sorted based on the number of data sets assigned to each.

Manage. Class Name

The name of the class assigned to the data sets in this group by the management class ACS routine.

Note that a summary line will appear for each unique combination of storage class, management class, data class, and storage group. The report groups are sorted based on the number of data sets assigned to each.

Storage Group Name

The name of the group assigned to the data sets in this group by the storage group ACS routine. Note that a summary line will appear for each unique combination of storage class, management class, data class, and storage group. The report groups are sorted based on the number of data sets assigned to each.

Data Class Name

The name of the class assigned to the data sets in this group by the data class ACS routine. Note that a summary line will appear for each unique combination of storage class, management class, data class, and storage group. The report groups are sorted based on the number of data sets assigned to each.

Number of Data Sets

A numeric value indicating the number of this owner's data sets assigned by the ACS routines to the classes and groups identified.

Percent of Owner Data Sets

The percentage of all data sets belonging to this owner that were assigned to this particular combination of classes and groups.

Space Allocated M-Bytes

The amount of space, reported in megabytes, that is being allocated by the data sets that were assigned to this particular combination of classes and groups. This value is calculated using data set allocation statistics extracted from CA MICS at the time the test cases were created.

Percent of Owner Megabytes

Based on the total megabytes allocated by all the data sets belonging to this owner, the percentage of that space allocated by the data sets assigned to this particular combination of classes and groups.

ACS TESTING RESULTS (STGELG)								
CA MICS StorageMate - OWNER								
ABC CORPORATION								
----- SYSID=SYSA Owner=APPL-DEVELOPMENT -----								
Storage Class Name	Manage. Class Name	Storage Group Name	Data Class Name		Number of Data Sets	Percent of Owner Data Sets		Space Allocated M-Bytes
								Percent of Owner Megabytes
STANDARD	LARGE	LARGE GP	DCPS		29	29.3		5239.8
STANDARD	STANDARD	TSOWRK	DCPO		27	27.3		43.5
STANDARD	NOMIGRT	TSOWRK	DCPS		12	12.1		1476.3
STANDARD	STANDARD	TSOWRK	(NONE)		12	12.1		2.6
STANDARD	STANDARD	TSOWRK	DCPS		10	10.1		0.9
STANDARD	STANDARD	TSOWRK	DCDA		6	6.1		46.4
STANDARD	LARGE	LARGE GP	(NONE)		1	1.0		25.4
VIOTEST		VIOTEST	DCPO		1	1.0		4.2
VIOTEST		VIOTEST	DCVS		1	1.0		39.7
-----					99			6878.9
OWNER								
----- SYSID=SYSA Owner=TECH-SUPPORT -----								
Storage Class Name	Manage. Class Name	Storage Group Name	Data Class Name		Number of Data Sets	Percent of Owner Data Sets		Space Allocated M-Bytes
								Percent of Owner Megabytes
STANDARD	LARGE	LARGE GP	DCPS		80	28.4		11296.9
STANDARD	STANDARD	TSOWRK	DCPS		62	22.0		21.0
STANDARD	STANDARD	TSOWRK	DCPO		47	16.7		41.7
STANDARD	STANDARD	TSOWRK	(NONE)		41	14.5		24.3
STANDARD	STANDARD	TSOWRK	DCVS		14	5.0		49.7
(NONE)			DCPS		7	2.5		236.5
STANDARD	STANDARD	TSOWRK	DCDA		7	2.5		79.5
STANDARD	STANDARD	COBGRP	DCPO		5	1.8		2.9
STANDARD	(NONE)	TSOWRK	DCPO		4	1.4		2.5
(NONE)			DCPO		3	1.1		1.4
STANDARD	(NONE)	TSOWRK	(NONE)		3	1.1		0.0
*RC=0008			DCPS		2	0.7		0.8
STANDARD	(NONE)	TSOWRK	DCPS		2	0.7		1.8
(NONE)			(NONE)		1	0.4		0.0
*RC=0008			DCPO		1	0.4		0.8
*RC=0012			DCDA		1	0.4		2.8
*RC=0012			DCPO		1	0.4		0.1
STANDARD	(NONE)	TSOWRK	DCVS		1	0.4		0.1
-----					282			11762.8
OWNER								

Figure 4-11. ACS Testing Results / Owner

Allocation Problems / Chart (STGECB)

Source: BATPGM file at the DETAIL timespan
BATJOB file at the DETAIL timespan
BAT_JS file at the DETAIL timespan

Function: This chart helps identify the amount of CPU time being expended in recovery activities for job abends related to DASD space shortages. These shortages can be caused by poor JCL parameters, by fragmented DASD space, or by a shortage of free space for certain departments or during certain time periods. This report gives the Storage Administrator a broad overview of the groups or users that are having problems, and helps identify specific time periods in which problems occur. This chart helps the Capacity Planner calculate the amount of processor resources wasted in recovery activities.

Features: The detection and calculation of resources expended doing recovery activities could save much time-consuming manual research. A wide variety of filtering and reporting options allow many views of the data, such as a broad system-wide view, a narrow view related to an individual time or department, or a view of a specific time period. This report could be run on a daily basis to give a general indication of storage system health. On days when anomalies occur, running the Allocations Problems analyses again with more detailed options could help isolate the problem. Both color graphics and printer graphics are supported by the report.

SMS

Issues: Users implementing SMS may want to identify departments having space-related abends as potential candidates for movement to SMS Storage Groups. After SMS is implemented, this report can help make sure abends are not occurring for managed users, and to alter SMS parameters if necessary.

COLUMN DESCRIPTIONS

Pie Slices

The pie slices represent the total amount of processor time (TCB plus SRB CPU time) required to

perform recovery for all the space-related failures within this group. This includes the CPU time of the failing steps, and the CPU time used by any steps running under the same job names during the recovery period. It does not include the CPU time used by the recovery steps. If no recovery steps are found for the jobs in this group, this will only include the CPU time used by the original failing steps.

Pie Labels

The labels on each slice identify the name of the group for which this summarization was performed. Depending on the values coded on the report option panel, this value can represent one or more of the BATPGM file accounting fields, or can be one of the following:

JOBCLASS = (Summary by Job Class of failing jobs)
ZONE = (Summary by ZONE of the failing jobs)
HOUR = (Summary by Hour of the failing jobs)
TYPE = (Summary by space-related abend type)




Figure 4-12. Allocation Problems / Chart

Allocation Problems / Summary (STGECB)

Source: BATPGM file at the DETAIL timespan
 BATJOB file at the DETAIL timespan
 BAT_JS file at the DETAIL timespan

Function: This report helps identify conditions that cause batch job abends related to DASD space shortages. These causes can include poor JCL parameters, fragmented DASD space, or a shortage of free space for certain departments or during certain time periods. This report helps the Storage Administrator focus on particular groups or users that are having problems, or to identify specific time periods in which problems are occurring. This report helps the Capacity Planner calculate the amount of human and machine resources wasted in recovery activities.

Features: The detection and calculation of resources expended doing recovery activities could save much time-consuming manual research. A wide variety of filtering and reporting options allow many views of the data, such as a broad system-wide view, a narrow view related to an individual time or department, or a view of a specific time period.

Generation of the summary report is automatic.

SMS

Issues: Users implementing SMS may want to identify departments having space-related abends as potential candidates for movement to SMS Storage Groups. After SMS is implemented, this report can be used to make sure abends are not occurring for managed users, and to alter SMS parameters if necessary.

COLUMN DESCRIPTIONS

Summary Group

The name of the group for which this summarization was performed. Depending on the values coded on the report option panel, this value can represent one or more of the BATPGM file accounting fields, or can be one of the following:

JOBCLASS = (Summary by Job Class of failing jobs)
 ZONE = (Summary by ZONE of the failing jobs)

HOURL = (Summary by Hour of the failing jobs)
TYPE = (Summary by space-related abend type)

Recovery Elapsed Time

The total amount of elapsed time necessary for the recovery process to complete for the jobs in this group. This is defined as the sum of the difference between the start time of the original failing step and the start time of the recovery step for all the jobs in this group. If no recovery steps were found for any of the jobs in this group, a value of zeros will be present.

Rerun Steps

The number of job steps (program executions) that had to occur to effect the recovery process for the failures within this group. This does not include the original failing steps, but does include the recovery steps, plus any steps that ran during the recovery periods with the same jobnames as the original failing jobs. This will have a value of zero if no recovery steps were found for any of the jobs in the group.

Rerun CPU Time

The amount of processor time (TCB plus SRB CPU time) required to perform recovery all the failures within this group. This includes the CPU time of the failing steps, and the CPU time used by any steps running under the same job names during the recovery period. It does not include the CPU time used by the recovery steps. If no recovery steps are found for the jobs in this group, this will only include the CPU time used by the original failing steps.

Rerun EXCPs

The number of I/O requests issued to perform recovery for this step. This includes the I/O counts of the failing step, and the I/O used by any steps running under the same job name during the recovery period. It does not include the I/O used by the recovery step. If no recovery step is found, this only includes the I/O counts used by the original failing step.

ALLOCATION PROBLEMS (STGECB)					2
CA MICS StorageMate - SUMMARY					
ABC CORPORATION					
----- SYSID=SYSA -----					
Summary Group	Recovery Elapsed Time	Rerun Steps	Rerun CPU Time	Rerun EXCPs	
SALES	0:03:51	2	0:00:01	541	
MARKETING	0:02:48	1	0:00:05	4936	
DEVELOPMENT	0:45:44	3	0:04:25	53584	
SUPPORT	1:00:21	2	0:05:48	77682	
TRAINING	0:04:35	1	0:00:10	2960	
-----	-----	-----	-----	-----	
SYSID	1:57:20	9	0:10:29	139703	
	=====	=====	=====	=====	
	1:57:20	9	0:10:29	139703	

Figure 4-13. Allocation Problems / Summary

Allocation Problems / Detail (STGECB)

Source: BATPGM file at the DETAIL timespan
BATJOB file at the DETAIL timespan
BAT_JS file at the DETAIL timespan

Function: This report helps to identify users or departments that are receiving batch job abends related to DASD space shortages. These problems can be caused by poor JCL parameters, fragmented DASD space, or by a shortage of free space. This report helps the Storage Administrator investigate space-related problems, and identify users that require JCL changes or more DASD space. This report helps the Capacity Planner calculate the amount of human and machine resources wasted in recovery activities.

Features: The left side of the report gives specific information about each abend, and is useful for doing detailed problem investigation. The right side of the report gives information about the recovery activities that occurred to correct the problem, which can save much time-consuming manual research. A wide variety of filtering and reporting options allow many views of the data, such as a broad system-wide view, or a narrow view related to an individual time or department.

SMS

Issues: Users implementing SMS may want to identify departments having space-related abends as potential candidates for movement to SMS Storage Groups. After SMS is implemented, this report can be used to make sure abends are not occurring for managed users, and to alter SMS parameters if necessary.

COLUMN DESCRIPTIONS

Jobname

The jobname of each batch job found within the input data that ended abnormally due to a space-related error.

Job Number/Class

The JES job number assigned to the failing job, and also the job class in which the job executed. One or both of these values may be blank if the

job number or class cannot be determined.

Step Start Date/Time

The date and time the failing step began.

Step Name

The step name from the EXEC JCL statement for the step that failed.

Program Name

The program name from the EXEC JCL statement that was executing when the failure occurred.

Step No.

The step number of the specified job within which the failure occurred.

Abend Code

The specific abort code that was received by the failing step. Refer to the IBM System Codes manual for an explanation of the meanings of each code.

Recovery Ended Time

The time at which recovery for the space-related error was completed. This is the time when the recovery step for the error started execution. The report defines a recovery step as a step having the same jobname, stepname, and program name as the originally abending step, which executes after the original failure within the time frame of the selected data cycles for this report. If a recovery step was not found, this time will be the same as the Step Start Date/Time.

Comments

This column will contain one of three values. If the value 'NO RECOVERY FOUND' is present, this means that no recovery step was found to recover from this error. The report defines a recovery step as a step having the same jobname, stepname, and program name as the originally abending step, which executes after the original failure. If the value 'SECONDARY ABEND' is present, it means that a recovery step was found, but that recovery step also received a space-related failure. Recovery for the original abend is considered complete when the secondary job begins, and recovery for the secondary abend will be reported in a future entry on the report. If the value 'RESTART STEP:name'

is present, this means that recovery was completed for this failure, and the stepname of the first step to run as part of the recovery process is listed in the comment.

Recovery Elapsed Time

The amount of time necessary for the recovery process to complete for this job. This is defined as the difference between the start time of the original failing step (Step Start Date/Time column) and the start time of the recovery step (Recovery Ended Time column). If no recovery step was found for this job, a value of zeros will be present. For jobs that experience a secondary abend during recovery, this will be the step start time for the recovery step that received the secondary abend.

Rerun Steps

The number of job steps (program executions) that had to occur to effect the recovery process for this job. This does not include the original step that received the abend, but does include the recovery step, plus any steps that ran during the recovery period with the same jobname as the original failing job. This will have a value of zero if no recovery step was found. When a secondary abend occurs, one will be added to this value to represent that step.

Rerun CPU Time

The amount of processor time (TCB plus SRB CPU time) required to perform recovery for this step. This includes the CPU time of the failing step, and the CPU time used by any steps running under the same job name during the recovery period. It does not include the CPU time used by the recovery step. If no recovery step is found, this only includes the CPU time used by the original failing step.

Rerun EXCPs

The number of I/O requests issued to perform recovery for this step. This includes the I/O counts of the failing step, and the I/O used by any steps running under the same job name during the recovery period. It does not include the I/O used by the recovery step. If no recovery step is found, this only includes the I/O counts used by the original failing step.

1

ALLOCATION PROBLEMS (STGECB)
CA MICS StorageMate - DETAIL
ABC CORPORATION

----- System Identifier=SYSA Jobclass=A -----

Jobname	Job Number /Class	Step Start Date/Time	Step Name	Program Name	Step No.	Abend Code	Recovery Ended Time	Comments	Recovery Elapsed Time	Rerun Steps	Rerun CPU Time	Rerun EXCPs
LNR231\$D	7231/A	18MAY07:10:36:00	MICS	SASLPA	11	SB37	11:27:41	RESTART STEP:MICS	0:51:41	1	0:02:57	38839
LNR298A	7394/A	18MAY07:16:04:22	C	IEV90	1	SE37	16:08:57	RESTART STEP:C	0:04:35	1	0:00:10	2960
LNR231\$D	7741/A	18MAY07:17:46:19	MICS	SASLPA	11	SB37	17:54:59	RESTART STEP:MICS	0:08:40	1	0:02:51	38843
-----									-----			
JOBCLASS									1:04:56	3	0:05:58	80642
N= 3												

----- System Identifier=SYS7 Jobclass=B -----

Jobname	Job Number /Class	Step Start Date/Time	Step Name	Program Name	Step No.	Abend Code	Recovery Ended Time	Comments	Recovery Elapsed Time	Rerun Steps	Rerun CPU Time	Rerun EXCPs
LNR118U1	7349/B	18MAY07:13:24:05	SAS	SASLPA	1	SB37	13:26:53	SECONDARY ABEND	0:02:48	1	0:00:03	2899
LNR118U1	7354/B	18MAY07:13:26:53	SAS	SASLPA	1	SB37	13:26:53	NO RECOVERY FOUND	0:00:00	0	0:00:02	2037
LNR050XZ	7554/B	18MAY07:15:46:11	COPY	IEBCOPY	1	SE37	15:46:52	SECONDARY ABEND	0:00:40	1	0:00:01	269
LNR050XZ	7556/B	18MAY07:15:46:52	COPY	IEBCOPY	1	SE37	15:50:03	RESTART STEP:COPY	0:03:11	1	0:00:01	272
LNR212CL	7636/B	18MAY07:16:53:22	SAS	SASLPA	1	SB37	17:39:06	RESTART STEP:SAS	0:45:44	3	0:04:25	53584
-----									-----			
JOBCLASS									0:52:24	6	0:04:31	59061
SYSID									1:57:20	9	0:10:29	139703
									=====	=====	=====	
									1:57:20	9	0:10:29	139703
N= 5												

Figure 4-14. Allocation Problems / Detail

ASTEX / Top Resource Consumers / Chart (STGEPF)

Source: ASTADS file at the DETAIL timespan

Function: This facility allows the storage administrator to quickly identify the primary users of resources critical to the operation of the storage I/O subsystem. Most performance analysts are familiar with the 80/20 rule, which states that 20% of your users are responsible for 80% of the resource usage on your system. Based on the increasing demands being made upon storage administrators, it is most effective to concentrate upon the usage of the system by its primary users. This facility allows the storage administrator to select one or more critical resource measures for which reporting will be done. The reports produced will quickly identify primary resource consumers, giving the storage administrator information about where tuning activities will produce the best results.

When potential problems have been detected by the summary reports or charts in this facility, a detail report option can be used to provide more information about the area in question. This will narrow specific problems down to the individual hour and data set level, if necessary.

Features: As stated above, a wide variety of data selection and summarization options are available. Data can be summarized using one or more nodes of the data set names, one or more CA MICS accounting fields, or using a user-defined table that relates data set prefixes with descriptions. Selection and reporting can also be done by date, hour, and zone. All of these options provide various views and levels of detail, making it easy to quickly identify and isolate problems.

The user is allowed to select the resource elements that can be used for reporting, as many different elements are available that measure storage system activity. Several of these may be selected in one report run, providing a large amount of information from a single report execution.

Output reports from this facility include tabular

reports, printer charts, and color graphic charts.

SMS

Issues: Users creating or modifying their SMS storage classes and groups must carefully monitor the I/O load on their existing applications, making sure that load is distributed so as to provide the desired performance for critical groups. These reports will help the user design storage classes and groups, and will allow the monitoring of resources after SMS implementation to make sure the SMS groups are still working correctly. This continued monitoring is essential, as even well-tuned groups may change as I/O workloads fluctuate.

ELEMENT DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Date

The date during which the measurements shown on this chart were taken. The interval at which measurements are taken, and the number of data cycles included in the report will affect the number of dates that appear in the output.

Hour

If summarization by hour was selected, the hour during which the measurements on the chart were taken will be displayed.

Zone

If summarization by zone was selected, the CA MICS zone during which the measurements on the chart were taken will be displayed.

Pie Chart Slices

The slices on the pie chart represent the amount of the selected storage resource being consumed by a particular group or user. The label on the slice identifies the group represented by the pie slice. Depending on the options selected, two numeric values may also appear on the chart associated with each slice. One represents the actual numeric value being represented by the

slice, while the other shows the percentage of the entire pie that slice represents. For example, assume that Total I/O Requests was the value being charted, and the numbers 5000 and 25% appeared for a given slice. That would mean that 5000 I/O requests were issued by this group, representing 25% of all I/O requests issued during this measurement interval.

Pie Chart Labels

The label associated with each pie slice indicates the group to which that slice applies. This may be assigned using one or more of the data set name nodes, one or more of the CA MICS accounting fields, or using a user-specified table that associates data set prefixes with descriptive labels.

Footnotes

The footnote at the bottom of each chart indicates the storage measurement represented by the chart. A separate chart will appear for each element selected in each time period selected. The possible values that are possible, and a short description of their meanings are as follows:

Total I/O Duration

The total duration of all I/O requests issued by this summary group during this measurement interval. This represents the sum of IOS queue time, pending time, disconnect time, and connect time for all requests in this group.

Total IOS Queue Time

The total amount of time spent by I/O requests in this summary group waiting on the IOS request queue to be scheduled.

Total I/O Pending Time

The total amount of time spent by I/O requests in this summary group in pending status. This represents the time an I/O request must wait for a clear path to the intended device.

Total I/O Disconnect Time

The total amount of time spent by I/O requests in this summary group in disconnect status. This represents the interval between the time the I/O request has been passed to the device

and the time the device is ready to transfer data.

Total I/O Connect Time

The total amount of time spent by I/O requests in this summary group in connect status. This represents the time that data is actually being transferred between the device and the processor.

Total Number of I/O Requests

The total number of I/O requests issued by all users in this summary group during this measurement interval.

Total RPS Time

The total amount of time spent by I/O requests in this summary group in RPS status. This represents the time the request is waiting while the DASD device is being positioned to the correct sector.

Total Seek I/O Requests

The total number of Seek I/O requests issued by users in this summary group during this measurement interval. A seek request causes the read/write head of the DASD device to position itself at a specified cylinder address.

Average Seek Distance in Cylinders

The average of all seek distances, expressed in cylinders, for each I/O request in this summary group for this measurement interval. This represents the average distance the DASD read/write heads had to move to satisfy the I/O requests for this group.

Average Response Time in Milliseconds

The average time, expressed in millionths of a second, it took for each I/O request to complete for the I/Os in this summary group during this measurement interval.

Total Cache Effective (Hit) Count

The total number of I/O requests for this summary group during this measurement interval that were processed by cache, and did not have to access the DASD device directly.

Total Cache Candidate I/O Requests

The total number of I/O requests for this summary group during this measurement interval that were considered to be cache candidates.

Total Non-Candidate I/O Requests

The total number of I/O requests for this summary group during this measurement interval that were not considered to be candidates for caching.

Total Tracks Loaded

The total number of tracks loaded into the cache on behalf of I/O requests in this summary group during this measurement interval.

Total DASD Fast Write Tracks Loaded

The total number of tracks loaded into the DASD Fast Write cache on behalf of I/O requests in this summary group during this measurement interval.

Total NVS Overflow Count

The total number of times that I/O requests by this summary group caused the non-volatile storage (NVS) cache area to become full and overflow during this measurement interval.

Total Dispatching Exceptions

The total number of dispatching exceptions that were caused by I/O requests in this summary group during this measurement interval. A dispatch exception is caused when excessive delay occurs getting an I/O dispatched.

Total Dsname Exceptions

The total number of data set name exceptions that were caused by I/O requests in this summary group during this measurement interval. A dsname exception is caused when excessive delay occurs at the data set level.

Total Path Exceptions

The total number of path exceptions that were caused by I/O requests in this summary group during this measurement interval. A path

exception is caused when excessive delay occurs at the I/O path level.

Total Seek Exceptions

The total number of seek exceptions that were caused by I/O requests in this summary group during this measurement interval.

Total Volume Exceptions

The total number of volume exceptions that were caused by I/O requests in this summary group during this measurement interval. A volume exception is caused when excessive delay occurs at the volume level.

Total Cross-Volume Exceptions

The total number of cross-volume exceptions that were caused by I/O requests in this summary group during this measurement interval. A cross-volume exception is caused when excessive delay occurs at the volume level, caused by activity from another system.

Total I/O When Volume Reserved

The total number of I/O requests issued by this summary group during this measurement interval when the intended device was reserved by another user.




Figure 4-15. ASTEX / Top Resource Consumers / Chart

ASTEX / Top Resource Consumers / Summary (STGEPF)

Source: ASTADS file at the DETAIL timespan

Function: This facility allows the storage administrator to quickly identify the primary users of resources critical to the operation of the storage I/O subsystem. Most performance analysts are familiar with the 80/20 rule, which states that 20% of your users are responsible for 80% of the resource usage on your system. Based on the increasing demands being made upon storage administrators, it is most effective to concentrate upon the usage of the system by its primary users. This facility allows the storage administrator to select one or more critical resource measures for which reporting will be done. The reports produced will quickly identify primary resource consumers, giving the storage administrator information about where tuning activities will produce the best results.

When potential problems have been detected by the summary reports or charts in this facility, a detail report option can be used to provide more information about the area in question. This will narrow specific problems down to the individual hour and data set level, if necessary.

Features: As stated above, a wide variety of data selection and summarization options are available. Data can be summarized using one or more nodes of the data set names, one or more of the CA MICS accounting fields, or using a user-defined table that relates data set prefixes with descriptions. Selection and reporting may also be done by date, hour, and zone. All of these options provide various views and levels of detail, making it easy to quickly identify and isolate problems.

The user is allowed to select the resource elements that can be used for reporting, because many different elements are available that measure storage system activity. Several of these may be selected in one report run, providing a large amount of information from a single report execution.

Output reports from this facility include tabular

reports, printer charts, and color graphic charts.

SMS

Issues: Users creating or modifying their SMS storage classes and groups must carefully monitor the I/O load on their existing applications, making sure that load is distributed so as to provide the desired performance for critical groups. These reports will help the user design storage classes and groups, and will allow the monitoring of resources after SMS implementation to make sure the SMS groups are still working correctly. This continued monitoring is essential, as even well-tuned groups may change as I/O workloads fluctuate.

COLUMN DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Date

The date during which the measurements shown on this report were taken. The interval at which measurements are taken, and the number of data cycles included in the report will affect the number of dates that appear in the output.

Hour

If summarization by hour was selected, the hour during which the measurements on the report were taken will be displayed.

Zone

If summarization by zone was selected, the CA MICS zone during which the measurements on the report were taken will be displayed.

Owner

Represents the group to which the displayed totals apply. This may be assigned using one or more of the data set name nodes, one or more of the CA MICS accounting fields, or using a user-specified table that associates data set prefixes with descriptive labels.

Data Columns

A separate data column will appear to the right of the Owner column. It will contain the measurement element(s) you have selected from the report options screen. If multiple elements are selected, each element will appear on a separate page. A maximum of 50 different owners will appear on each page, representing the 50 largest users of the selected element. The report will be sorted so that the largest users of the resource will appear first. The label associated with each possible data column, and a short description of what each value represents, appears below:

Total I/O Duration

The total duration of all I/O requests issued by this summary group during this measurement interval. This represents the sum of IOS queue time, pending time, disconnect time, and connect time for all requests in this group.

Total IOS Queue Time

The total amount of time spent by I/O requests in this summary group waiting on the IOS request queue to be scheduled.

Total Pending Time

The total amount of time spent by I/O requests in this summary group in pending status. This represents the time an I/O request must wait for a clear path to the intended device.

Total Disconnect Time

The total amount of time spent by I/O requests in this summary group in disconnect status. This represents the interval between the time the I/O request has been passed to the device and the time the device is ready to transfer data.

Total Connect Time

The total amount of time spent by I/O requests in this summary group in connect status. This represents the time that data is actually being transferred between the device and the processor.

Total I/O Requests

The total number of I/O requests issued by all users in this summary group during this measurement interval.

Total RPS Time

The total amount of time spent by I/O requests in this summary group in RPS status. This represents the time the request is waiting while the DASD device is being positioned to the correct sector.

Total Seek I/O Requests

The total number of seek I/O requests issued by users in this summary group during this measurement interval. A seek request causes the read/write head of the DASD device to position itself at a specified cylinder address.

Average Seek Size in CYLS

The average seek distance, expressed in cylinders, for each I/O request in this summary group for this measurement interval. This represents the average distance the DASD read/write heads had to move to satisfy the I/O requests for this group.

Average Response Time in Milliseconds

The average time, expressed in millionths of a second, it took for each I/O request to complete for the I/Os in this summary group during this measurement interval.

Total Cache Effective (Hit) Count

The total number of I/O requests for this summary group during this measurement interval that were processed by cache, and did not have to access the DASD device directly.

Total Cache Candidate I/O Requests

The total number of I/O requests for this summary group during this measurement interval that were considered to be cache candidates.

Total Non-Candidate I/O Requests

The total number of I/O requests for this summary group during this measurement interval that were not considered to be candidates for caching.

Total Tracks Loaded

The total number of tracks loaded into the cache on behalf of I/O requests in this summary group during this measurement interval.

Total DFW Tracks Loaded

The total number of tracks loaded into the DASD

Fast Write cache on behalf of I/O requests in this summary group during this measurement interval.

Total NVS Overflow Count

The total number of times that I/O requests by this summary group caused the non-volatile storage (NVS) cache area to become full and overflow during this measurement interval.

Total Dispatching Exceptions

The total number of dispatching exceptions that were caused by I/O requests in this summary group during this measurement interval. A display exception is caused when excessive delay occurs getting an I/O dispatched.

Total Dsname Exceptions

The total number of data set name exceptions that were caused by I/O requests in this summary group during this measurement interval. A dsname exception is caused when excessive delay occurs at the data set level.

Total Path Exceptions

The total number of path exceptions that were caused by I/O requests in this summary group during this measurement interval. A path exception is caused when excessive delay occurs at the I/O path level.

Total Seek Exceptions

The total number of seek exceptions that were caused by I/O requests in this summary group during this measurement interval.

Total Volume Exceptions

The total number of volume exceptions that were caused by I/O requests in this summary group during this measurement interval. A volume exception is caused when excessive delay occurs at the volume level.

Total Cross-Volume Exceptions

The total number of cross-volume exceptions that were caused by I/O requests in this summary group during this measurement interval. A cross-volume exception is caused when excessive delay occurs at the volume level, caused by activity from another system.

Total I/O When Volume Reserved

The total number of I/O requests issued by this summary group during this measurement interval when the intended device was reserved by another user.

ASTEX / TOP RESOURCE CONSUMERS (STGEPPF)	
CA MICS StorageMate	
ABC CORPORATION	
----- SYSID=SYSA Date=05APR91 Hour=9 -----	
Owner	Total I/O Requests
SUPPORT	204548
SALES	181906
PAYROLL	151365
FINANCE	75403
DATA-BASE-PROD	40616
SYS-PROGS	10119
CORPORATE	7598
MARKETING	125
DATA-BASE-TEST	7
-----	-----
HOUR	671687

Figure 4-16. ASTEX / Top Resource Consumers / Summary

ASTEX / Top Resource Consumers / Detail (STGEPPF)

Source: ASTADS file at the DETAIL timespan

Function: This facility allows the storage administrator to quickly identify the primary users of resources critical to the operation of the storage I/O subsystem. Most performance analysts are familiar with the 80/20 rule, which states that 20% of your users are responsible for 80% of the resource usage on your system. Based on the increasing demands being made upon storage administrators, it is most effective to concentrate upon the usage of the system by its primary users. This facility allows the storage administrator to select one or more critical resource measures for which reporting will be done. The reports produced will quickly identify primary resource consumers, giving the storage administrator information about where tuning activities will produce the best results.

When potential problems have been detected by the summary reports or charts in this facility, a detail report option can be used to provide more information about the area in question. This will narrow specific problems down to the individual hour and data set level, if necessary.

Features: As stated above, a wide variety of data selection and summarization options are available. Data can be summarized using one or more nodes of the data set names, one or more of the CA MICS accounting fields, or using a user-defined table that relates data set prefixes with descriptions. Selection and reporting can also be done by date, hour, and zone. All of these options provide various views and levels of detail, making it easy to quickly identify and isolate problems.

The user is allowed to select the resource elements that may be used for reporting, because many different elements are available that measure storage system activity. Several of these can be selected in one report run, providing a large amount of information from a single report execution.

Output reports from this facility include tabular reports, printer charts, and color graphic charts.

SMS

Issues: Users creating or modifying their SMS storage classes and groups must carefully monitor the I/O load on their existing applications, making sure that load is distributed so as to provide the desired performance for critical groups. These reports will help the user design storage classes and groups, and will allow the monitoring of resources after SMS implementation to make sure the SMS groups are still working correctly. This continued monitoring is essential, as even well-tuned groups may change as I/O workloads fluctuate.

COLUMN DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Owner

Represents the group to which the displayed totals apply. This can be assigned using one or more of the data set name nodes, one or more of the CA MICS accounting fields, or using a user-specified table that associates data set prefixes with descriptive labels.

Date

The date during which the measurements shown on this report were taken. The interval at which measurements are taken, and the number of data cycles included in the report will affect the number of dates that appear in the output.

Hour

If summarization by hour was selected, the hour during which the measurements on the report were taken will be displayed.

Zone

If summarization by zone was selected, the CA MICS zone during which the measurements on the report were taken will be displayed.

Time

The time stamp taken from the detail record for this particular group of totals. This represents the approximate ending time of the measurement interval.

Data Set Name

The name of the data set for which this line of measurements apply.

Volume Serial Number

The volume serial number of the volume on which the indicated data set resides.

Data Columns

One or more data columns will appear between the Time and the Data Set Name columns. These will contain the measurement element(s) you have selected from the report options screen. If you select more than one element with the same alphabetic character, all these elements will appear on the same detail report. The label associated with each possible data column, and a short description of what each value represents, appears below:

Total I/O Duration

The total duration of all I/O requests issued by this summary group during this measurement interval. This represents the sum of IOS queue time, pending time, disconnect time, and connect time for all requests in this group.

Total IOS Queue Time

The total amount of time spent by I/O requests in this summary group waiting on the IOS request queue to be scheduled.

Total Pending Time

The total amount of time spent in pending status by I/O requests in this summary group. This represents the time an I/O request must wait for a clear path to the intended device.

Total Disconnect Time

The total amount of time spent in disconnect status by I/O requests in this summary group. This represents the interval between the time the I/O request has been passed to the device and the time the device is ready to transfer data.

Total Connect Time

The total amount of time spent in connect status by I/O requests in this summary group. This represents the time that data is actually being transferred between the device and the processor.

Total I/O Requests

The total number of I/O requests issued by all users in this summary group during this measurement interval.

Total RPS Time

The total amount of time spent by I/O requests in this summary group in RPS status. This represents the time the request is waiting while the DASD device is being positioned to the correct sector.

Total Seek I/O Requests The total number of Seek I/O requests issued by users in this summary group during this measurement interval. A seek request causes the read/write head of the DASD device to position itself at a specified cylinder address.

Average Seek Size in CYLS

The average seek distance, expressed in cylinders, for each I/O request in this summary group for this measurement interval. This represents the average distance the DASD read/write heads had to move to satisfy the I/O requests for this group.

Average Response Time in Milliseconds

The average time, expressed in millionths of a second, it took for each I/O request to complete for the I/Os in this summary group during this measurement interval.

Total Cache Effective (Hit) Count

The total number of I/O requests for this summary group during this measurement interval that were processed by cache, and did not have to access the DASD device directly.

Total Cache Candidate I/O Requests

The total number of I/O requests for this summary group during this measurement interval that were considered to be cache candidates.

Total Non-Candidate I/O Requests

The total number of I/O requests for this summary group during this measurement interval that were considered to not be candidates for caching.

Total Tracks Loaded

The total number of tracks loaded into the cache on behalf of I/O requests in this summary group during this measurement interval.

Total DFW Tracks Loaded

The total number of tracks loaded into the DASD Fast Write cache on behalf of I/O requests in this summary group during this measurement interval.

Total NVS Overflow Count

The total number of times that I/O requests by this summary group caused the non-volatile storage (NVS) cache area to become full and overflow during this measurement interval.

Total Dispatching Exceptions

The total number of dispatching exceptions that were caused by I/O requests in this summary group during this measurement interval. A dispatch exception is caused when excessive delay occurs getting an I/O dispatched.

Total Dsname Exceptions

The total number of data set name exceptions that were caused by I/O requests in this summary group during this measurement interval. A dsname exception is caused when excessive delay occurs at the data set level.

Total Path Exceptions

The total number of path exceptions that were caused by I/O requests in this summary group during this measurement interval. A path exception is caused when excessive delay occurs at the I/O path level.

Total Seek Exceptions

The total number of seek exceptions that were caused by I/O requests in this summary group during this measurement interval.

Total Volume Exceptions

The total number of volume exceptions that were caused by I/O requests in this summary group during this measurement interval. A volume exception is caused when excessive delay occurs at the volume level.

Total Cross-Volume Exceptions

The total number of cross-volume exceptions that were caused by I/O requests in this summary group during this measurement interval. A cross-volume exception is caused when excessive delay occurs at the volume level, caused by activity from another system.

Total I/O When Volume Reserved

The total number of I/O requests issued by this summary group during this measurement interval when the intended device was reserved by another user.

ASTEX / TOP RESOURCE CONSUMERS (STGEPPF)				
CA MICS StorageMate - DETAIL				
ABC CORPORATION				
----- System Identifier=SYSA Owner=CORPORATE Date=05APR91 Hour=9 -----				
Time	Total Connect Time	Total I/O Requests	Data Set Name	Volume Serial Number
9:00	0:00:00.4	33	FIN.TEST.MASTER.LIB	FIN008
9:00	0:00:00.5	143	FIN1.PROCLIB	GLMRS1
9:00	0:00:04.3	1681	GENTEST.DEBUG11.FILE	TS0059
9:00	0:00:01.6	315	GEN.GLSYSP.GENMENU	GLMRS1
9:00	0:00:00.4	50	FIN.SPF.FINPLIB	SYSWXA
9:00	0:00:02.0	457	FIN.GLSYST.FINLOAD	SYSRS1
9:00	0:00:02.7	832	FIN.GLSYSP.FINMENU	GLMRS1
9:00	0:00:02.0	459	FIN.GLSYSP.FINPENU	GLMRS1
9:00	0:00:00.4	76	FIN999.TEST01.CONLIB1	TS0026
9:30	0:00:00.2	44	FIN.MIGR.TESTFILE	FIN008
9:30	0:00:01.0	245	COR.MSTR99.ISFPLIB	GLMRS1
9:30	0:00:01.5	265	GEN.GLSYSP.GENMENU	GLMRS1
9:30	0:00:00.5	81	FIN.SPF.FINPLIB	SYSWXA
9:30	0:00:01.7	391	FIN.GLSYST.FINLOAD	SYSRS1
9:30	0:00:03.5	992	FIN.GLSYSP.FINMENU	GLMRS1
9:30	0:00:02.8	626	FIN.GLSYSP.FINPENU	GLMRS1
9:30	0:00:00.1	40	FIN999.TEST01.CONLIB1	TS0065
-----	-----	-----		
HOURL	0:00:25.6	6730		
	=====	=====		

Figure 4-17. ASTEX / Top Resource Consumers / Detail

Daily Application Resource Usage / Plot (STGERG)

Source: VCADAA file at the DAYS timespan
VCA_VS file at the DAYS timespan
HSM MIG file at the DAYS timespan
HSBBAC file at the DAYS timespan

Function: This facility allows the storage administrator to monitor key storage indicators on a day-by-day basis. These indicators enable the user to project short-term storage needs, both at the application level and the system level. The indicators reported by this facility are also useful for detecting problems with storage resources that are not being used efficiently. These reports, along with the Monthly Application Resource Usage reports, provide both a short view and a long view of storage capacity and performance problems.

The ability to process data from multiple sources makes the reports produced even more valuable. Options allow the user to target a single storage media, such as real DASD or HSM backup storage, or to group data from multiple media on the same chart. This second type of reporting is useful when analyzing the flow of data across multiple types of storage media.

The types of data that can be selected for reporting include the number of data sets, amount of space being allocated, percentages of data sets and space in certain categories, and the amount of allocated versus used space.

Features: Summarization options allow you to analyze the data based on its owner (using one or more of the CA MICS accounting fields), its SMS storage class, its SMS management class, or its location (such as real DASD, or HSM level 1 migration).

Data filtering allows you to include data from selected summary groups. For example, you could choose to report on specific departments, groups or individuals, or on a specific SMS class.

Output reports from this facility include tabular reports, printer plots, or color graphic plots. When plots are requested, up to nine different

elements can be plotted on the same chart.

SMS

Issues: The ability to summarize by an SMS storage or management class makes it easy to monitor the short-term activity within each of your defined classes. Many of the elements you can select for reporting are useful for monitoring an SMS migration effort, or to make sure SMS is operating correctly after migration is complete.

ELEMENT DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Media

Storage Class

Management Class

Group

The top line of each page also indicates the summary group to which the values on this plot apply. Depending on the report option chosen, this summarization will be done by media (DASD, VSAM, MIGRATE, or BACKUP), SMS storage class, SMS management class, or based on the CA MICS accounting codes in the data. The label to the left of the summary value will be one of the values shown above, depending on the summarization option chosen.

Vertical Axis (Value)

The numeric values distributed vertically down the left side of each page represent the scale of the numeric values being plotted on that page. When requesting color graphics plots, you can specify the values that will appear on this axis. This should be done with care, because all plots and groups will have to adapt to the values you specify. For printer plots, or when values are not specified, SAS will generate the values automatically, based on the content of the plotted data. We suggest that these SAS defaults be used whenever possible, as they will avoid many potential problems. Care should also be taken when specifying multiple elements on one plot, that the elements chosen will have a similar range

of values.

Horizontal Axis (Date)

Plots on the horizontal axis represent the date corresponding to the value being plotted. Each date is presented in the format 'ddmmmy'. The number of entries appearing on the horizontal axis is determined by the number of cycles of CA MICS data specified as input to the inquiry.

Footnotes

The footnote(s) section of each graph is important, as it indicates the names of the data element(s) that are plotted on the graph. Up to nine elements can be plotted on the same graph, in which case 9 footnotes will appear at the bottom of each page. For printer plots, an alphabetic character A-I will appear to the left of each description, identifying the character used on the plot to represent that particular element. On color graphic plots, the letter does not appear, as the color of the footnote should correspond with the color of the plot point/line on the graph. If different colors do not appear for each point and footnote, your MICS options should be modified to specify different colors. A brief description of the possible values in the footnote section, and what those values represent, is given below:

Data Set Count: Total

Represents the total number of data sets found in the input data for this particular summary group.

Data Set Count: SMS-Managed

Represents the number of data sets found in the input data for this particular summary group that are managed by SMS.

Data Set Count: Unmanaged

Represents the number of data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Data Sets: Managed

Represents the percentage of all the data sets found in the input data for this particular summary group that are managed by

SMS.

Percentage of Data Sets: Unmanaged

Represents the percentage of all the data sets found in the input data for this particular summary group that are not managed by SMS.

Space Allocated (MB): Total

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group.

Space Allocated (MB): Managed

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are managed by SMS.

Space Allocated (MB): Unmanaged

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Space: Managed

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are managed by SMS.

Percentage of Space: Unmanaged

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not managed by SMS.

Data Set Count: Non-VSAM

Represents the number of data sets found in the input data for this particular summary group that are not VSAM data sets.

Data Set Count: VSAM

Represents the number of data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Data Sets: Non-VSAM

Represents the percentage of data sets found

in the input data for this particular summary group that are not VSAM data sets.

Percentage of Data Sets: VSAM

Represents the percentage of data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Space Used: All Data

Represents the percentage of space allocated versus used for all the data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the data sets in this group by the total tracks allocated. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Managed

Represents the percentage of space allocated versus used for the SMS-managed data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the managed data sets in this group by the total tracks allocated by the managed data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Unmanaged

Represents the percentage of space allocated versus used for the unmanaged data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the unmanaged data sets in this group by the total tracks allocated by the unmanaged data sets. Note that HSM data sets are not included in this calculation.

Space Allocated (MB): Non-VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not VSAM data sets.

Space Allocated (MB): VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the

input data for this particular summary group that are VSAM data sets.

Percentage of Space: Non-VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not VSAM data sets.

Percentage of Space: VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are VSAM data sets.

Percentage of Space Used: Non-VSAM

Represents the percentage of space allocated versus used for the non-VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the non-VSAM data sets in this group by the total tracks allocated by the non-VSAM data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: VSAM

Represents the percentage of space allocated versus used for the VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the VSAM data sets in this group by the total tracks allocated by the VSAM data sets. Note that HSM data sets are not included in this calculation.

Data Set Count: HSM Level 1

Represents the number of data sets found in the input data for this particular summary group that reside on DFHSM migration level 1 storage.

Data Set Count: HSM Level 2

Represents the number of data sets found in the input data for this particular summary group that reside on DFHSM migration level 2 storage.

Data Set Count: HSM Backup

Represents the number of data sets found in the input data for this particular summary group that reside on DFHSM backup storage. Note that multiple backup copies of the same data set do not effect this count, nor is there any attempt made to indicate which of these data sets may also reside on primary storage.

Space Allocated (MB): HSM Level 1

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on DFHSM migration level 1 storage.

Space Allocated (MB): HSM Level 2

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on DFHSM migration level 2 storage.

Space Allocated (MB): HSM Backup

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on DFHSM backup storage. This total includes space used for multiple backup copies of the same data set.

Percentage of Data Sets: HSM Level 1

Represents the percentage of data sets found in the input data for this particular summary group that reside on DFHSM migration level 1 storage.

Percentage of Data Sets: HSM Level 2

Represents the percentage of data sets found in the input data for this particular summary group that reside on DFHSM migration level 2 storage.

Percentage of Data Sets: HSM Backup

Represents the percentage of data sets found in the input data for this particular summary group that reside on DFHSM backup storage. Note that multiple backup copies of the same data set do not affect this count, nor is any

attempt made to indicate which of these data sets may also reside on primary storage.

Percentage of Space: HSM Level 1

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on DFHSM migration level 1 storage.

Percentage of Space: HSM Level 2

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on DFHSM migration level 2 storage.

Percentage of Space: HSM Backup

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on DFHSM backup storage. This total includes space used for multiple backup copies of the same data set.




Figure 4-18. Daily Application Resource Usage / Plot

Daily Application Resource Usage / Summary (STGERG)

Source: VCADAA file at the DAYS timespan
VCA_VS file at the DAYS timespan
HSMIMG file at the DAYS timespan
HSBBAC file at the DAYS timespan

Function: This facility allows the storage administrator to monitor key storage indicators on a day-by-day basis. These indicators enable the user to project short-term storage needs, both at the application level and the system level. The indicators reported by this facility are also useful for detecting trends that could indicate potential problems or areas in which storage resources are not being used efficiently. These reports, along with the Monthly Application Resource Usage reports, provide both a short view and a long view of storage capacity and performance problems.

The ability to process data from multiple sources makes the reports produced even more valuable. Options allow the user to target a single storage media, such as real DASD or HSM backup storage, or to group data from multiple media on the same chart. This second type of reporting is useful when analyzing the flow of data across multiple types of storage media.

The types of data that can be selected for reporting include the number of data sets, amount of space being allocated, percentages of data sets and space in certain categories, and the amount of allocated versus used space.

Features: Summarization options allow you to analyze the data based on its owner (using one or more of the CA MICS accounting fields), its SMS storage class, its SMS management class, or its location (such as real DASD or DFHSM level 1 migration).

Data filtering allows you to include data from selected summary groups. For example, you could choose to report on specific departments groups or individuals, or on a specific SMS class.

Output reports from this facility include tabular reports, printer plots, or color graphic plots.

When plots are requested, up to nine different elements can be plotted on the same chart.

SMS

Issues: The ability to summarize by an SMS storage or management class makes it easy to monitor the long-term activity within each of your defined classes. Many of the elements you can select for reporting are useful for monitoring an SMS migration effort, or to make sure SMS is operating correctly after migration is complete.

COLUMN DESCRIPTIONS

TITLE 4

The fourth report title line on each tabular report page indicates the indicator or element that is being reported on that page. Note the values displayed for each item will vary, depending on the input sources specified for the report. Any or all of the following values may appear:

Data Set Count:

Total Represents the total number of data sets found in the input data for this particular summary group.

Data Set Count: SMS-Managed

Represents the number of data sets found in the input data for this particular summary group that are managed by SMS.

Data Set Count: Unmanaged

Represents the number of data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Data Sets: Managed

Represents the percentage of all the data sets found in the input data for this particular summary group that are managed by SMS.

Percentage of Data Sets: Unmanaged

Represents the percentage of all the data sets found in the input data for this particular summary group that are not managed by SMS.

Space Allocated (MB): Total

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group.

Space Allocated (MB): Managed

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are managed by SMS.

Space Allocated (MB): Unmanaged

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Space: Managed

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are managed by SMS.

Percentage of Space: Unmanaged

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not managed by SMS.

Data Set Count: Non-VSAM

Represents the number of data sets found in the input data for this particular summary group that are not VSAM data sets.

Data Set Count: VSAM

Represents the number of data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Data Sets: Non-VSAM

Represents the percentage of data sets found in the input data for this particular summary group that are not VSAM data sets.

Percentage of Data Sets: VSAM

Represents the percentage of data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Space Used: All Data

Represents the percentage of space allocated versus used for all the data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the data sets in this group by the total tracks allocated. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Managed

Represents the percentage of space allocated versus used for the SMS-managed data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the managed data sets in this group by the total tracks allocated by the managed data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Unmanaged

Represents the percentage of space allocated versus used for the unmanaged data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the unmanaged data sets in this group by the total tracks allocated by the unmanaged data sets. Note that HSM data sets are not included in this calculation.

Space Allocated (MB): Non-VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not VSAM data sets.

Space Allocated (MB): VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Space: Non-VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not VSAM data sets.

Percentage of Space: VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage

of that space belonging to data sets that are VSAM data sets.

Percentage of Space Used: Non-VSAM

Represents the percentage of space allocated versus used for the non-VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the non-VSAM data sets in this group by the total tracks allocated by the non-VSAM data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: VSAM

Represents the percentage of space allocated versus used for the VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the VSAM data sets in this group by the total tracks allocated by the VSAM data sets. Note that DFHSM data sets are not included in this calculation, because they are always fully used.

Data Set Count: HSM Level 1

Represents the number of data sets found in the input data for this particular summary group that reside on DFHSM migration level 1 storage.

Data Set Count: HSM Level 2

Represents the number of data sets found in the input data for this particular summary group that reside on DFHSM migration level 2 storage.

Data Set Count: HSM Backup

Represents the number of data sets found in the input data for this particular summary group that reside on DFHSM backup storage. Note that multiple backup copies of the same data set do not effect this count, nor is there any attempt made to indicate which of these data sets may also reside on primary storage.

Space Allocated (MB): HSM Level 1

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on DFHSM migration level 1 storage.

Space Allocated (MB): HSM Level 2

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on DFHSM migration level 2 storage.

Space Allocated (MB): HSM Backup

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on HSM backup storage. This total includes space used for multiple backup copies of the same data set.

Percentage of Data Sets: HSM Level 1

Represents the percentage of data sets found in the input data for this particular summary group that reside on DFHSM migration level 1 storage.

Percentage of Data Sets: HSM Level 2

Represents the percentage of data sets found in the input data for this particular summary group that reside on DFHSM migration level 2 storage.

Percentage of Data Sets: HSM Backup

Represents the percentage of data sets found in the input data for this particular summary group that reside on DFHSM backup storage. Note that multiple backup copies of the same data set do not effect this count, nor is there any attempt made to indicate which of these data sets may also reside on primary storage.

Percentage of Space: HSM Level 1

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on DFHSM migration level 1 storage.

Percentage of Space: HSM Level 2

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on DFHSM migration level 2 storage.

Percentage of Space: HSM Backup

Based on the total space allocated by all the

data sets in this summary group, the percentage of that space belonging to data sets that reside on DFHSM backup storage. This total includes space used for multiple backup copies of the same data set.

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Media**Storage Class****Management Class****Owner**

The first column of the report indicates the summary group to which the totals on each line apply. Depending on the report option chosen, this summarization will be done by media (DASD, VSAM, MIGRATE, or BACKUP), SMS storage class, SMS management class, or based on the CA MICS accounting codes in the data. The label over the column will be one of the values shown above, depending on the summarization option chosen.

Value ddmmmyy

Indicates the value for this particular element for this particular summary group during the date indicated in the column heading. The date is represented in the format 'ddmmmyy'.

Compound Growth Rate (%)

Reports the daily compound growth rate (CGR) calculated for this group using the values from the days shown on this report. A value of 5.00, for example, indicates a growth rate of approximately 5% per day for this group, based on the daily values included in this report.

DAILY APPLICATION RESOURCE USAGE (STGERG)							
CA MICS StorageMate							
ABC CORPORATION							
Data Set Count: Total							
----- SYSID=SYSA -----							
Owner	Value 07APR91	Value 08APR91	Value 09APR91	Value 10APR91	Value 11APR91	Value 12APR91	Compound Growth Rate (%)
DEPT:100	654	657	774	791	751	764	3.16
DEPT:110	1565	1589	1613	1606	1625	1614	0.62
DEPT:120	1438	1535	1577	1644	1684	1696	3.36
DEPT:125A	9117	9257	9409	9430	9435	9518	0.86
DEPT:125B	522	522	522	521	521	521	-0.04
DEPT:130	495	366	366	366	365	365	-5.91
DEPT:140	56	56	56	56	56	56	0.00
DEPT:145	805	623	450	753	796	882	1.84
DEPT:150	1261	1265	1266	1274	1273	1269	0.13
DEPT:157	64	63	64	62	62	61	-0.96
DEPT:160	106	107	107	108	121	126	3.52
DEPT:170	739	747	760	761	769	774	0.93
DEPT:200C	565	572	617	638	640	637	2.43
DEPT:210	426	433	443	439	437	436	0.47
DEPT:220	932	933	936	939	909	924	-0.17
DEPT:230	1375	1399	1342	1298	1355	1494	1.67
DEPT:240	82	82	82	81	82	81	-0.25
DEPT:270	2778	2841	2909	3223	2844	2806	0.20
DEPT:300	595	644	662	692	676	671	2.43
DEPT:350	70	70	70	70	70	75	1.39
DEPT:700A	29	29	29	29	29	29	0.00
DEPT:999	11	11	11	9	9	9	-3.93
-----	-----	-----	-----	-----	-----	-----	-----
SYSID	23685	23801	24065	24790	24509	24808	
	=====	=====	=====	=====	=====	=====	

Figure 4-19. Daily Application Resource Usage / Summary

Daily ASTEX Cache Statistics / Plot (STGEPE)

Source: ASTAVO file at the DAYS timespan
 HARDVA file at the DAYS timespan

Function: This facility allows the storage administrator to monitor key storage indicators related to the operation of DASD cache control units. While DASD cache has become an important part of the total storage hierarchy, the lack of convenient and powerful reporting tools has made it difficult to manage this resource. The introduction of ASTEX data into CA MICS should help correct this problem, because it gives storage administrators the ability to quickly and easily monitor this important resource.

Cache-related problems are normally caused by too many users needing cache resources at the same time. Detection and correction of these problems requires the ability to analyze the data at varying levels of granularity. This facility allows selection and summarization by specific dates, hours, and CA MICS zones. Selection can also be limited to the data for a specific volume, SMS storage group, ASTEX contention group, or MVS logical control unit. This type of multi-dimensional analysis can detect problems that might otherwise remain hidden.

Features: As stated above, a wide variety of options are available in terms of data selection and summarization. These options make it easy to run several reports at varying levels of detail to quickly identify and isolate a problem.

Output reports from this facility include tabular reports, printer plots, or color graphic plots. When plots are requested, up to nine different elements may be plotted on the same chart.

SMS

Issues: The ability to select and summarize by an SMS storage group allows you to monitor and tune your cache resources and your SMS groups so they deliver optimum performance. The ability to summarize the report using a Volume Group Table allows you to simulate the operation of alternate storage group definitions. Thus, you can easily

move volumes between storage groups, monitor the results, and then finally make the change in your group definitions when the results are acceptable.

ELEMENT DESCRIPTIONS

SYSID/S

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies. The shorter value is used with color graphic plots.

Volume/Vol

Storage Group/Grp

Logical Ctl Unit/LCUID

Contention Group/CGrp

The top line of each page also indicates the summary group to which the values on this plot apply. Depending on the report option chosen, this summarization will be done by volume, SMS storage group, MVS logical control unit, or ASTEX contention group. The shorter value listed above is used with color graphic plots.

Date

When summarization by hour or zone is requested, the date will appear at the top of each chart to indicate the date for which the plotted data applies.

Vertical Axis (Value)

The numeric values distributed vertically down the left side of each page represent the scale of the numeric values being plotted on that page. When requesting color graphics plots, you can specify the values that will appear on this axis. This should be done with care, because all plots and groups will have to adapt to the values you specify. For printer plots, or when values are not specified, SAS will generate the values automatically, based on the content of the plotted data. We suggest that these SAS defaults be used whenever possible, because they will avoid many potential problems. Care should also be taken when specifying multiple elements on one plot, that the elements chosen will have a similar range of values.

Horizontal Axis (Date/Zone/Hour)

Plots on the horizontal axis represent the date, hour, or zone corresponding to the value being plotted. Report options are used to determine which of these three values will be used. When date summarization is specified, the date appears in the format 'ddmmyy', and the number of dates plotted is determined by the number of cycles of CA MICS data specified as input to the report.

Footnotes

The footnote(s) section of each graph is important, as it indicates the names of the data element(s) that are plotted on the graph. Up to 9 elements can be plotted on the same graph, in which case 9 footnotes will appear at the bottom of each page. For printer plots, an alphabetic character A-I will appear to the left of each description, identifying the character used on the plot to represent that particular element. On color graphic plots, the letter does not appear, because the color of the footnote should correspond with the color of the plot point/line on the graph. If different colors do not appear for each point and footnote, your MICF options should be modified to specify different colors. A brief description of the possible values in the footnote section, and what those values represent, is given below:

I/O Requests: Total

The total number of I/O requests received for this group during this measurement interval.

I/O Requests: Cache Candidate

The number of I/O requests received that were considered to be eligible for caching.

I/O Requests: Read The total number of read I/O requests that were considered to be eligible for caching.

I/O Requests: Read Sequential

The total number of read sequential I/O requests that were considered to be eligible for caching.

I/O Requests: DASD Fast Write

The number of I/O requests received that were considered candidates for DASD Fast Write

caching.

I/O Requests: Cache Fast Write

The Number of I/O requests received that were considered candidates for Cache Fast Write caching.

I/O Requests: Non-Cache Candidate

The total number of I/O requests received that were not considered to be cache candidates.

I/O Requests: Write

The total number of I/O requests received that were not considered to be cache candidates because they were write requests.

I/O Requests: Bypass Cache

The total number of I/O requests received that were not considered to be cache candidates because they requested no caching.

I/O Requests: IOS Inhibited

The total number of I/O requests received that were not considered to be cache candidates because IOS had inhibited them due to cache activity.

I/O Requests: ASTEX Inhibited

The total number of I/O requests received that were not considered to be cache candidates because ASTEX had inhibited them due to cache activity.

I/O Requests: Other

The total number of I/O requests received that were not considered to be cache candidates for reasons other than those listed above.

Cache Effective: Total

The total number of I/O requests that were successfully processed by cache (cache hits).

Cache Effective: Read

The total number of read I/O requests that were successfully processed by cache (cache hits).

Cache Effective: Read Sequential

The total number of read sequential I/O

requests that were successfully processed by cache (cache hits).

Cache Effective: DASD Fast Write

The total number of I/O requests that were successfully processed by DASD Fast Write cache.

Cache Effective: Cache Fast Write

The total number of I/O requests that were successfully processed by Cache Fast Write cache.

Total Tracks Loaded into Cache

The total number of tracks loaded into cache during this measurement interval.

Total DASD Fast Write Track Loads

The total number of tracks loaded into DASD Fast Write cache during this measurement interval.

NVS Overflow Percentage

The percentage of times that requests for non-volatile storage exceeded NVS capacity and caused an overflow condition.

Average Demotion Time

The average amount of time a storage page stayed in cache memory, before being replaced by a more active page.

Percent Effective: All I/O

The average percentage of all I/O requests received that were handled by cache storage.

Percent Effective: Cache Candidates

The average percentage of all cache candidate I/O requests received that were handled by cache storage.

Percent Effective: Read

The average percentage of all read I/O requests received that were handled by cache storage.

Percent Effective: Read Sequential

The average percentage of all read sequential I/O requests received that were handled by

cache storage.

Percent Effective: DASD Fast Write

The average percentage of all DASD Fast Write requests received that were handled by DASD Fast Write.

Percent Effective: Cache Fast Write

The average percentage of all Cache Fast Write requests received that were handled by Cache Fast Write.

Percent Demotion Time < 30 seconds

The percentage of time during this interval that the average demotion time was less than 30 seconds. Demotion time indicates the amount of time a page can remain unreferenced before being removed from cache.

Percent Demotion Time < 1 minute

The percentage of time during this interval that the average demotion time was less than 1 minute. Demotion time indicates the amount of time a page can remain unreferenced before being removed from cache.

Percent Demotion Time < 2 minutes

The percentage of time during this interval that the average demotion time was less than two minutes. Demotion time indicates the amount of time a page can remain unreferenced before being removed from cache.

Percent Demotion Time > 2 minutes

The percentage of time during this interval that the average demotion time was greater than two minutes. Demotion time indicates the amount of time a page can remain unreferenced before being removed from cache.




Figure 4-20. Daily ASTEX Cache Statistics / Plot

Daily ASTEX Cache Statistics / Summary (STGEPE)

Source: ASTAVO file at the DAYS timespan
 HARDVA file at the DAYS timespan

Function: This facility allows the Storage Administrator to monitor key storage indicators related to the operation of DASD cache control units. While DASD cache has become an important part of the total storage hierarchy, the lack of convenient and powerful reporting tools has made it difficult to manage this resource. The introduction of ASTEX data into CA MICS should help correct this problem, because it gives Storage Administrators the ability to quickly and easily monitor this important resource.

Cache-related problems are normally caused by too many users needing cache resources at the same time. Detection and correction of these problems requires the ability to analyze the data at varying levels of granularity. This facility allows selection and summarization by specific dates, hours, and CA MICS zones. Selection can also be limited to the data for a specific volume, SMS storage group, ASTEX contention group, or MVS logical control unit. This type of multi-dimensional analysis can detect problems that might otherwise remain hidden.

Features: As stated above, a wide variety of options are available in terms of data selection and summarization. These options make it easy to run several reports at varying levels of detail to quickly identify and isolate a problem.

Output reports from this facility include tabular reports, printer plots, and color graphic plots. When plots are requested, up to nine different elements can be plotted on the same chart.

SMS

Issues: The ability to select and summarize by an SMS storage group allows you to monitor and tune your cache resources and your SMS groups so they deliver optimum performance. The ability to summarize the report using a Volume Group Table allows you to simulate the operation of alternate storage group definitions. Thus, you can easily

move volumes between storage groups, monitor the results, and then finally make the change in your group definitions when the results are acceptable.

COLUMN DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Volume

Storage Group

Logical Ctl Unit

Contention Group

The top line of each page also indicates the summary group to which the values on this report apply. Depending on the report option chosen, this summarization will be done by volume, SMS storage group, MVS logical control unit, or ASTEX contention group.

Date

Indicates the date for which the reported data applies. Depending on the summarization option chosen, the date will appear either as part of the heading or as the first column of data.

Hour

When summarization by hour is requested, this indicates the hour to which the reported data applies.

Zone

When summarization by zone is requested, this indicates the CA MICS zone to which the reported data applies.

Other Data Columns

Other printed columns that appear on the report correspond to the elements you have selected to monitor. All of the elements you select from the report options screen with the same letter will appear on the same report. A brief description of the possible columns that may appear here, and what those columns represent, is given below:

Total I/O Requests

The total number of I/O requests received for

this group during this measurement interval.

Candidate I/O Requests

The number of I/O requests received that were considered to be eligible for caching.

Read I/O Requests

The total number of read I/O requests that were considered to be eligible for caching.

Read Seq I/O Requests

The total number of read sequential I/O requests that were considered to be eligible for caching.

DASD FW I/O Requests

The number of I/O requests received that were considered candidates for DASD Fast Write caching.

Cache FW I/O Requests

The Number of I/O requests received that were considered candidates for Cache Fast Write caching.

Non-Cand I/O Requests

The total number of I/O requests received that were not considered to be cache candidates.

Write I/O Requests

The total number of I/O requests received that were not considered to be cache candidates because they were write requests.

Cache Byp I/O Requests

The total number of I/O requests received that were not considered to be cache candidates because they requested no caching.

IOS Inhib I/O Requests

The total number of I/O requests received that were not considered to be cache candidates because IOS had inhibited them due to cache activity.

ASTEX Inhib I/O Requests

The total number of I/O requests received that were not considered to be cache

candidates because ASTEX had inhibited them due to cache activity.

Other I/O Requests

The total number of I/O requests received that were not considered to be cache candidates for reasons other than those listed above.

Cache Hits Total

The total number of I/O requests that were successfully processed by cache (cache hits).

Cache Hits Read

The total number of read I/O requests that were successfully processed by cache (cache hits).

Cache Hits Rd Seq

The total number of read sequential I/O requests that were successfully processed by cache (cache hits).

Cache Hits DFW

The total number of I/O requests that were successfully processed by DASD Fast Write cache.

Cache Hits CFW

The total number of I/O requests that were successfully processed by Cache Fast Write cache.

Total Tracks Loaded

The total number of tracks loaded into cache during this measurement interval.

DFW Tracks Loaded

The total number of tracks loaded into DASD Fast Write cache during this measurement interval.

NVS Overflow Percentage

The percentage of times that requests for non-volatile storage exceeded NVS capacity and caused an overflow condition.

Average Demotion Time

The average amount of time a storage page

stayed in cache memory, before being replaced by a more active page.

% Hits All I/O

The average percentage of all I/O requests received that were handled by cache storage.

% Hits Cand I/O

The average percentage of all cache candidate I/O requests received that were handled by cache storage.

% Hits Read I/O

The average percentage of all read I/O requests received that were handled by cache storage.

% Hits Rd Seq I/O

The average percentage of all read sequential I/O requests received that were handled by cache storage.

% Hits DFW I/O

The average percentage of all DASD Fast Write requests received that were handled by DASD Fast Write.

% Hits CFW I/O

The average percentage of all Cache Fast Write requests received that were handled by Cache Fast Write.

% Demote Time < 30 sec

The percentage of time during this interval that the average demotion time was less than 30 seconds. Demotion time indicates the amount of time a page can remain unreferenced before being removed from cache.

% Demote Time < 1 min

The percentage of time during this interval that the average demotion time was less than one minute. Demotion time indicates the amount of time a page can remain unreferenced before being removed from cache.

% Demote Time < 2 min

The percentage of time during this interval that the average demotion time was less than

two minutes. Demotion time indicates the amount of time a page can remain unreferenced before being removed from cache.

% Demote Time > 2 min

The percentage of time during this interval that the average demotion time was greater than two minutes. Demotion time indicates the amount of time a page can remain unreferenced before being removed from cache.

DAILY ASTEX CACHE STATISTICS (STGEPE)						
CA MICS StorageMate						
ABC CORPORATION						
----- SYSID=SYSA Storage Group=TSOLRG Date=05APR91 -----						
Hour	Non-Cand I/O Requests	Write I/O Requests	Cache Byp I/O Requests	IOS Inhib I/O Requests	ASTEX Inhib I/O Requests	Other I/O Requests
0	4709	935	0	3042	257	475
1	53338	16948	0	32549	1450	2391
2	24238	10268	0	10446	1324	2200
3	18880	14069	0	102	3851	858
4	3568	704	0	18	2528	318
5	1000	607	0	38	61	294
6	1902	722	0	10	816	354
7	2077	686	0	36	1026	329
8	7931	4727	0	233	649	2322
9	3465	961	0	761	1303	440
17	2444	902	0	699	358	485
18	1573	754	0	188	271	360
19	2239	1367	0	22	526	324
20	1886	1473	0	6	96	311
21	1047	619	0	20	95	313
22	1041	610	0	0	131	300
23	1048	620	0	0	128	300
-----	-----	-----	-----	-----	-----	-----
DATE	132386	56972	0	48170	14870	12374

Figure 4-21. Daily ASTEX Cache Statistics / Summary

Daily DASD Resource Usage / Plot (STGEJE)

Source: VCAVOA file at the DETAIL timespan

Function: This facility allows the storage administrator to monitor key storage indicators on a day-by-day basis. These indicators enable the user to project short-term storage needs, both at the hardware unit level (volume, volume group, or device type) and the system level. The indicators reported by this facility are also useful for detecting problems with storage resources that are not being used efficiently. These reports, along with the tasks such as Monthly DASD Resource Usage reports, provide both a short view and a long view of storage capacity and performance problems, measured at the hardware level.

The Daily/Monthly DASD Resource Usage reports complement the Daily/Monthly Application Resource Usage reports to give a complete picture of storage usage and capacity. While the application reports tend to measure storage capacity and consumption at the application level, the DASD reports take a more hardware-oriented look at similar data. These reports are useful for tasks such as planning for new DASD or balancing volumes between volume groups or SMS storage groups. Both sets of reports allow you to measure and control growth at the application and hardware levels.

The types of data that can be selected for reporting include the number of data sets, the amount of space being allocated, percentages of data sets and space in certain categories, the amount of allocated versus used space, and other measures of performance and capacity.

Features: Summarization options allow you to analyze the data based on volume serial number, device type, volume group or SMS storage group, or to provide totals for each system.

Data filtering allows you to include data from selected groups. For example, you can choose to report on specific volumes, device types, or SMS storage groups.

Output reports from this facility include tabular

reports, printer plots, or color graphic plots. When plots are requested, up to nine different elements can be plotted on the same chart.

Many users choose to run the CA MICS DASD Collector multiple times during the day and then combine those multiple measurements into the same cycle of CA MICS data. The reports in this facility will detect this condition and alter their output to reflect this "detail mode" of operation. The values on each report and graph will then reflect individual measurements for each day rather than just one value per day. Detail mode will be activated when 3 or more unique measurements are found for the same volume on any one day. It is suggested you do not combine detailed and normal data cycles in the same report execution.

SMS

Issues: The ability to summarize by SMS storage groups makes it easy to monitor the short-term activity within each of your defined groups. Many of the elements you can select for reporting are useful for monitoring an SMS migration effort, or to make sure SMS is operating correctly after migration is complete. Using a Volume Group Table, alternate storage group configurations may be designed and simulated before being implemented.

ELEMENT DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Date

When detailed reporting is active, the date will appear at the top of each page in the format 'ddmmyy'. All data on the remainder of the page will apply to this date.

Volume

Volume Group

Device Type

Unless summarization by SYSID only is chosen, the descriptive line at the top of each page will also indicate the summary group to which the values on

this plot apply. Depending on the report option chosen, this summarization will be done by volume, volume group (SMS storage group), or device type. The label to the left of the summary value will be one of the values shown above, depending on the summarization option chosen.

Vertical Axis (Value)

The numeric values distributed vertically down the left side of each page represent the scale of the numeric values being plotted on that page. When requesting color graphics plots, you can specify the values that will appear on this axis. This should be done with care, because all plots and groups will have to adapt to the values you specify. For printer plots, or when values are not specified, SAS will generate the values automatically, based on the content of the plotted data. We suggest that these SAS defaults be used whenever possible, as they will avoid many potential problems. Care should also be taken when specifying multiple elements on one plot, that the elements chosen will have a similar range of values.

Horizontal Axis (Date)

Plots on the horizontal axis represent the date corresponding to the value being plotted. Each date is presented in the format 'ddmmmyy'. The number of entries appearing on the horizontal axis is determined by the number of cycles of CA MICS data specified as input to the inquiry.

Horizontal Axis (Time)

When detailed reporting is active, the plots on the horizontal axis represent the measurement time corresponding to the value being plotted. Each time is presented in the format 'hh:mm:ss'. The number of entries appearing on the horizontal axis is determined by the number of different measurements found each particular date.

Footnotes

The footnotes section of each graph is important, as it indicates the names of the data element(s) that are plotted on the graph. Up to nine elements can be plotted on the same graph, in which case nine footnotes will appear at the bottom of each page. For printer plots, an

alphabetic character ranging from A through I will appear to the left of each description, identifying the character used on the plot to represent that particular element. On color graphic plots, the letter does not appear, as the color of the footnote should correspond with the color of the plot point or line on the graph. If different colors do not appear for each point and footnote, modify your MICF options to specify different colors. A brief description of the possible values in the footnote section, and what those values represent, is given below:

Data Set Count: Total

Represents the total number of data sets found in the input data for this particular summary group.

Data Set Count: SMS-Managed

Represents the number of data sets found in the input data for this particular summary group that are managed by SMS.

Data Set Count: Unmanaged

Represents the number of data sets found in the input data for this particular summary group that are not managed by SMS.

Data Set Count: Non-VSAM

Represents the number of data sets found in the input data for this particular summary group that are not VSAM data sets.

Data Set Count: VSAM

Represents the number of data sets found in the input data for this particular summary group that are VSAM data sets.

Space Allocated (MB): Total

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group.

Space Allocated (MB): Managed

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are managed by SMS.

Space Allocated (MB): Unmanaged

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not managed by SMS.

Space Allocated (MB): Non-VSAM

Represents the total space allocated, in megabytes, by all the non-VSAM data sets found in the input data for this particular summary group.

Space Allocated (MB): VSAM

Represents the total space allocated, in megabytes, by all the VSAM data sets found in the input data for this particular summary group.

Percentage of Data Sets: Managed

Represents the percentage of all the data sets found in the input data for this particular summary group that are managed by SMS.

Percentage of Data Sets: Unmanaged

Represents the percentage of all the data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Data Sets: Non-VSAM

Represents the percentage of all the data sets found in the input data for this particular summary group that are non-VSAM data sets.

Percentage of Data Sets: VSAM

Represents the percentage of all the data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Space: Managed

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are managed by SMS.

Percentage of Space: Unmanaged

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not managed by SMS.

Percentage of Space: Non-VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not VSAM.

Percentage of Space: VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are VSAM.

Percentage of Space Used: All Data

Represents the percentage of space allocated versus used for all the data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the data sets in this group by the total tracks allocated. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Non-VSAM

Represents the percentage of space allocated versus used for the non-VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the non-VSAM data sets in this group by the total tracks allocated by the non-VSAM data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: VSAM

Represents the percentage of space allocated versus used for the VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the VSAM data sets in this group by the total tracks allocated by the VSAM data sets. Note that HSM data sets are not included in this calculation.

Total Storage Capacity (MB)

Represents the total published storage capacity for all the volumes in this particular summary group, expressed in megabytes.

Total Capacity Allocated (MB)

Represents the total amount of space within this particular summary group that has been allocated, and is no longer available for allocation by other data sets. This value is expressed in megabytes, and represents the space allocated by data sets and system areas such as VTOCs.

Total Capacity Free (MB)

Represents the total amount of space within this particular summary group that is available for allocation by other data sets. This value is expressed in megabytes.

Percentage of Capacity Allocated

Represents the percentage of space within this particular summary group that has been allocated, and is no longer available for allocation by other data sets. This value represents the space allocated by data sets and system areas such as VTOCs.

Percentage of Capacity Free

Represents the percentage of space within this particular summary group that is available for allocation by other data sets.

Average Fragmentation Index

Represents the average of all the fragmentation indexes for all the volumes in this summary group. The fragmentation index is used as a measurement of space fragmentation on each volume, with a low number indicating little fragmentation, and higher numbers indicating greater amounts of fragmentation.

Free VTOC Index Records

Represents the total number of free VTOC index records for all of the volumes in this particular summary group. For volumes with indexed VTOCs, this number indicates the amount of index expansion space available.




Figure 4-22. Daily DASD Resource Usage / Plot

Daily DASD Resource Usage / Summary (STGEJE)

Source: VCAVOA file at the DETAIL timespan

Function: This facility allows the storage administrator to monitor key storage indicators on a day-by-day basis. These indicators enable the user to project short-term storage needs, both at the hardware unit level (volume, volume group, or device type) and the system level. The indicators reported by this facility are also useful for detecting problems with storage resources that are not being used efficiently. These reports, along with the Monthly DASD Resource Usage reports, provide both a short view and a long view of storage capacity and performance problems, measured at the hardware level.

The Daily/Monthly DASD Resource Usage reports complement the Daily/Monthly Application Resource Usage reports to give a complete picture of storage usage and capacity. While the application reports tend to measure storage capacity and consumption at the application level, the DASD reports take a more hardware-oriented look at similar data. These reports are useful for tasks such as planning for new DASD, or balancing volumes between volume groups or SMS storage groups. Both sets of reports allow you to measure and control growth at the application and the hardware levels.

The types of data that can be selected for reporting include the number of data sets, the amount of space being allocated, percentages of data sets and space in certain categories, the amount of allocated versus used space, and other measures of performance and capacity.

Features: Summarization options allow you to analyze the data based on volume serial number, device type, volume group or SMS storage group, or to provide totals for each system.

Data filtering allows you to include data from selected groups. For example, you can choose to report on specific volumes, device types, or SMS storage groups.

Output reports from this facility include tabular reports, printer plots, or color graphic plots. When plots are requested, up to nine different elements can be plotted on the same chart.

Many users choose to run the CA MICS DASD Collector multiple times during the day and then combine those multiple measurements into the same cycle of CA MICS data. The reports in this facility will detect this condition and alter their output to reflect this "detail mode" of operation. The values on each report and graph will then reflect individual measurements for each day rather than just one value per day. Detail mode will be activated when 3 or more unique measurements are found for the same volume on any one day. It is suggested you do not combine detailed and normal data cycles in the same report execution.

SMS

Issues: The ability to summarize by SMS storage groups makes it easy to monitor the short-term activity within each of your defined groups. Many of the elements you can select for reporting are useful for monitoring an SMS migration effort, or to make sure SMS is operating correctly after migration is complete. Using a Volume Group Table, alternate storage group configurations may be designed and simulated before being implemented.

COLUMN DESCRIPTIONS

TITLE 4

The fourth report title line on each tabular report page indicates the indicator or element that is being reported on that page. Note the values displayed for each item will depend on the input sources specified for the report. Any or all of the following values may appear:

Data Set Count: Total

Represents the total number of data sets found the input data for this particular summary group.

Data Set Count: SMS-Managed

Represents the number of data sets found in the input data for this particular summary group

that are managed by SMS.

Data Set Count: Unmanaged

Represents the number of data sets found in the input data for this particular summary group that are not managed by SMS.

Data Set Count: Non-VSAM

Represents the number of data sets found in the input data for this particular summary group that are not VSAM data sets.

Data Set Count: VSAM

Represents the number of data sets found in the input data for this particular summary group that are VSAM data sets.

Space Allocated (MB): Total

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group.

Space Allocated (MB): Managed

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are managed by SMS.

Space Allocated (MB): Unmanaged

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not managed by SMS.

Space Allocated (MB): Non-VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not VSAM data sets.

Space Allocated (MB): VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Data Sets: Managed

Represents the percentage of all the data sets found in the input data for this particular

summary group that are managed by SMS.

Percentage of Data Sets: Unmanaged

Represents the percentage of all the data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Data Sets: Non-VSAM

Represents the percentage of all the data sets found in the input data for this particular summary group that are not VSAM.

Percentage of Data Sets: VSAM

Represents the percentage of all the data sets found in the input data for this particular summary group that are VSAM.

Percentage of Space: Managed

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are managed by SMS.

Percentage of Space: Unmanaged

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not managed by SMS.

Percentage of Space: Non-VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not VSAM data sets.

Percentage of Space: VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are VSAM data sets.

Percentage of Space Used: All Data

Represents the percentage of space allocated versus used for all the data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the data sets in this group by the total tracks allocated. Note that HSM data sets are not included in this

calculation.

Percentage of Space Used: Non-VSAM

Represents the percentage of space allocated versus used for the non-VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the non-VSAM data sets in this group by the total tracks allocated by the non-VSAM data sets. Note that HSM data sets are not included in this calculation.

Represents the percentage of space allocated versus used for the VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the VSAM data sets in this group by the total tracks allocated by the VSAM data sets. Note that DFHSM data sets are not included in this calculation, because they are always fully used.

Total Storage Capacity (MB)

Represents the total published storage capacity for all the volumes in this particular summary group, expressed in megabytes.

Total Capacity Allocated (MB)

Represents the total amount of space within this particular summary group that has been allocated, and is no longer available for allocation by other data sets. This value is expressed in megabytes, and represents the space allocated by data sets and system areas such as VTOCs.

Total Capacity Free (MB)

Represents the total amount of space within this particular summary group that is available for allocation by other data sets. This value is expressed in megabytes.

Percentage of Capacity Allocated

Represents the percentage of space within this particular summary group that has been allocated, and is no longer available for allocation by other data sets. This value

represents the space allocated by data sets and system areas such as VTOCs.

Percentage of Capacity Free

Represents the percentage of space within this particular summary group that is available for allocation by other data sets.

Average Fragmentation Index

Represents the average of all the fragmentation indexes for all the volumes in this summary group. The fragmentation index is used as a measurement of space fragmentation on each volume, with a low number indicating little fragmentation, and higher numbers indicating greater amounts of fragmentation.

Free VTOC Index Records

Represents the total number of free VTOC index records for all of the volumes in this particular summary group. For volumes with indexed VTOCs, this number indicates the amount of index expansion space available.

SYSID

Unless summarization by SYSID is chosen, this value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Date

When detailed reporting is active, the date will appear at the top of each page in the format 'ddmmyy'. All data on the remainder of the page will apply to this date.

SYSID

Volume

Volume Group

Device Type

The left-most data column on each report line indicates the summary group to which the values on that line apply. Depending on the report option chosen, this summarization will be done by volume, volume group (SMS storage group), device type, or system (SYSID). The label at the top of the column will be one of the values shown above, depending on the summarization option chosen.

Value ddmmmyy

Indicates the value for this particular element for this particular summary group during the date indicated in the column heading. The date is represented in the format 'ddmmmyy'.

Value hh:mm:ss

When detailed reporting is active, this indicates the value for this particular element for this particular summary group at the measurement time indicated in the column heading. The time is represented in the format 'hh:mm:ss'.

Compound Growth Rate (%)

Reports the compound growth rate (CGR) calculated for this group using the values from the days/times shown on this report. A value of 5.00, for example, indicates a growth rate of approximately 5% per day/interval for this group, based on the values included in the report.

DAILY DASD RESOURCE USAGE (STGEJE)						
CA MICS StorageMate						
ABC CORPORATION						
Space Allocated (MB): Total						
----- SYSID=5008 -----						
Volume Group	Value 01JAN92	Value 02JAN92	Value 03JAN92	Value 04JAN92	Value 05JAN92	Compound Growth Rate (%)
LARGE	20878	21654	23508	.	22893	2.33
TECHSUP	11349	11435	11416	.	11491	0.31
TEMP	304	408	214	.	2111	62.29
TSO	31911	31765	32908	.	33262	1.04

SYSID	64442	65261	68047	0	69757	
	=====	=====	=====	=====	=====	
	64442	65261	68047	0	69757	

Figure 4-23. Daily DASD Resource Usage / Summary

Daily HSM DASD Usage / Chart (STGEJC)

Source: HSMVOL file at the DETAIL timespan
VCAVOA file at the DETAIL timespan

Function: This facility reports the effectiveness of HSM in managing the DASD volumes under its control. The reports produced will enable you to determine when HSM parameters need to be adjusted, and when volume pool changes need to be made. Up to 10 different reports can be generated from this facility, and each report may be generated in tabular or graphic format. Color graphics are also supported. These reports monitor such items as the results of daily space management, the excess data occupancy that exists on your volumes, the frequency and success rate for interval migration, and other usage statistics for your DASD and HSM resources.

The data used for these reports is gathered directly from HSM by the daily run of the CA MICS DASD Space Collector. It represents a view of your storage resource as seen by HSM and, as such, may have results that are slightly different from other StorageMate reports. Be advised that HSM will only generate statistics for a volume if HSM activity occurred that day for that volume. This means that the amount of data produced for low activity days will be less than for busy days. Keep this in mind when viewing such metrics as total storage capacity, as these values may tend to change slightly from day to day. Some reports may also show no activity because of the values of your HSM options. For example, the values for interval migration will always be zero if you have specified no interval migration.

Features: Depending on the reports selected, most of the HSM statistics can be summarized by volume, HSM volume type (primary or migration), or volume group (SMS storage group). Because storage group is not kept in the HSM data, these values are extracted from the VCAVOA file and merged with the HSM data to create the final report.

All reports can be produced as tabular reports, printer graphics reports or color graphics reports.

SMS

Issues: Users planning for SMS could use these reports to tune HSM and determine the optimum HSM thresholds when defining SMS management classes. After SMS implementation has begun, the reports should be used on an ongoing basis to make sure adequate space exists in each storage group, and to make sure HSM is managing each storage group as designed.

ELEMENT DESCRIPTIONS

Chart types

Up to four different types of charts can be generated by this facility:

Migration Effectiveness

This chart reports the effectiveness of HSM daily space management activities. Different chart versions can be produced to report by volume, volume type, and/or volume group.

Excess Occupancy Analysis

This chart reports the amount of space occupied by data sets that did not migrate, but could migrate if necessary. Different chart versions can be produced to report by volume, volume type, and/or volume group.

Interval Migration Activity

This chart reports the number of times HSM started interval migration, and the number of times it was effective. Different chart versions can be produced to report by volume, volume type, or volume group.

Total Capacity Analysis

This chart shows capacity and usage values on a daily basis. It can be used to monitor the capacity and occupancy of all the volumes known to HSM.

While only one chart example is shown in this section, all four chart types will be described.

Common Elements

The following elements are common to most of the charts produced:

SYSID

The System Identifier to which this data applies.

Volume Serial Number (Volume)**HSM Volume Type (Level)****Volume Group (VolGrp)**

Identifies the type of summarization requested. This represents summarization based on individual volume, HSM volume type (0=primary or 1=migration) or volume group (SMS storage group). The alternate values shown in parentheses appear when using color graphics.

Date

Identifies the date associated with the values on the chart, shown as 'ddmmmyy'.

Migration Effectiveness

The following descriptions apply to the charts produced to show HSM Migration Effectiveness:

Occupancy Before Daily Space Management

Represents the percentage of the volume or volume group that was allocated before HSM daily space management began.

Occupancy After Daily Space Management

Represents the percentage of the volume or volume group that was allocated after HSM daily space management completed. The differences in these two values shows the percentage of the volume or group freed by HSM.

Excess Occupancy Analysis

The following descriptions apply to the charts produced to analyze Excess Occupancy:

Storage Capacity in Megabytes

Represents the total storage capacity of the volume or volume group expressed in megabytes.

Allocated Space in Megabytes

Represents the amount of space on the volume or volume group that was allocated after HSM

daily space management completed. The difference between the Capacity value and the Allocated value represents free space.

Excess Occupancy Space in Megabytes

Represents the amount of space, in megabytes, allocated by data sets that HSM did not migrate, but could have migrated if more free space were still required. As this value moves closer to zero, the ability for HSM to guarantee free space on this volume or group is reduced.

Interval Migration Activity

The following descriptions apply to the charts produced to analyze Interval Migration:

Total Interval Migrations

Represents the number of times HSM started interval migrations for this volume or volume group. Interval migrations are started at a user-specified interval if the free space on a volume has dropped below a user-specified value.

Failed Interval Migrations

Represents the number of times that HSM interval migration was started, but the interval migration could not free enough space to meet a user-specified value. Excessive numbers of interval migrations and failed migrations indicate a potential space shortage or a threshold value in that may be incorrect.

Total Capacity Analysis

The following descriptions apply to the charts produced to analyze Total Capacity:

Total Capacity (Primary and Level 1)

Represents the total storage capacity of all the volumes known to HSM, whether those volumes are primary or Level 1 migration volumes. This value is reported in megabytes.

Capacity of Primary Volumes

Represents the total storage capacity of all the primary volumes known to HSM, reported in

megabytes. The difference between this value and the Total Capacity value would indicate the capacity of the HSM Level 1 migration volumes.

Total Allocated (Primary and Level 1)

Represents the amount of space, in megabytes, allocated by data sets that reside on either primary volumes or HSM Level 1 migration volumes.

Allocated on Primary Volumes Only

Represents the allocated space for those data sets residing only on primary volumes known to HSM, reported in megabytes. The difference between this value and the Total Allocated value indicate the space allocated by data sets on HSM Level 1 migration volumes.




Figure 4-24. Daily HSM DASD Usage / Chart

Daily HSM DASD Usage / Summary (STGEJC)

Source: HSMVOL file at the DETAIL timespan
VCAVOA file at the DETAIL timespan

Function: This facility reports the effectiveness of HSM in managing the DASD volumes under its control. The reports produced will enable you to determine when HSM parameters need to be adjusted, and when volume pool changes need to be made. Up to 10 different reports can be generated from this facility, and each report may be generated in tabular or graphic format. Color graphics are also supported. These reports monitor such items as the results of daily space management, the excess data occupancy that exists on your volumes, the frequency and success rate for interval migration, and other usage statistics for your DASD and HSM resources.

The data used for these reports is gathered directly from HSM by the daily run of the CA MICS DASD Space Collector. It represents a view of your storage resource as seen by HSM and, as such, may have results that are slightly different from other StorageMate reports. Be advised that HSM will only generate statistics for a volume if HSM activity occurred that day for that volume. This means that the amount of data produced for low activity days will be less than for busy days. Keep this in mind when viewing such metrics as total storage capacity, as these values may tend to change slightly from day to day. Some reports may also show no activity because of the values of your HSM options. For example, the values for interval migration will always be zero if you have specified no interval migration.

Features: Depending on the reports selected, most of the HSM statistics can be summarized by volume, HSM volume type (primary or migration), or volume group (SMS storage group). Because storage group is not kept in the HSM data, these values are extracted from the VCAVOA file and merged with the HSM data to create the final report.

All reports can be produced as tabular reports, printer graphics reports or color graphics reports.

SMS

Issues: Users planning for SMS could use these reports to tune HSM and determine the optimum HSM thresholds when defining SMS management classes. After SMS implementation has begun, the reports should be used on an ongoing basis to make sure adequate space exists in each storage group, and to make sure HSM is managing each storage group as designed.

COLUMN DESCRIPTIONS

Report types

Up to four different types of reports can be generated by this facility:

Migration Effectiveness

This report shows the effectiveness of HSM daily space management activities. Different report versions can be produced to summarize the data by volume, volume type, and volume group.

Excess Occupancy Analysis

This report monitors the amount of space occupied by data sets that did not migrate, but could migrate if necessary. Different report versions can be produced to summarize the data by volume, volume type, and volume group.

Interval Migration Activity

This report shows the number of times HSM started interval migration, and the number of times it was effective. Different report versions can be produced to summarize the data by volume, volume type, and/or volume group.

Total Capacity Analysis

This report shows capacity and usage values on a daily basis. It can be used to monitor the capacity and occupancy of all the volumes known to HSM.

While only one report example is shown in this section, all four report types will be described.

Common Elements

The following elements are common to most of the reports produced:

SYSID

The System Identifier to which this data applies.

Volume Serial Number

HSM Volume Type

Volume Group

Identifies the type of summarization requested. This represents summarization based on individual volume, HSM volume type (0=primary or 1=migration) or volume group (SMS storage group).

Date

The left-most column of each report will list the date associated with the values shown on that report line. The date is shown in the format 'ddmmyy'.

Migration Effectiveness

The following descriptions apply to the reports produced to show HSM Migration Effectiveness:

Occupancy % Before Migration

Represents the percentage of the volume or volume group that was allocated before HSM daily space management began.

Occupancy % After Migration

Represents the percentage of the volume or volume group that was allocated after HSM daily space management completed.

Percentage Freed

Represents the percentage of the volume or group freed by HSM daily space management.

Excess Occupancy Analysis

The following descriptions apply to the reports produced to analyze Excess Occupancy:

Percentage Excess Occupancy

Represents the percentage of the volume or group allocated by data sets that HSM did not migrate, but could have migrated if more free space were still required. Large values here

are good, as they represent excess space that could be freed if needed.

Percentage Not Eligible

Represents the percentage of the volume or group allocated by data sets that HSM cannot migrate, as they don't meet the migration thresholds. Small values here are good, as they represent only a small number of data sets that HSM could not move if necessary.

Percentage Unallocated

Represents the percentage of the volume or group currently available for new data set allocations. Large values here are good, as they represent unused capacity. The sum of these three percentage columns should equal 100%.

Total Capacity (MB)

The total storage capacity of this volume or volume group, expressed in megabytes.

Interval Migration Activity

The following descriptions apply to the reports produced to analyze Interval Migration:

Total Interval Migrations

Represents the number of times HSM started interval migrations for this volume or volume group. Interval migrations are started at a user-specified interval if the free space on a volume has dropped below a user-specified value.

Successful Interval Migrations

Represents the number of times that HSM interval migration was successful, meaning that it was able to free enough space on the volume or group to meet a user-specified threshold.

Percentage Migrations Successful

The percentage of interval migrations for this volume or group that were successful. This is calculated using the values in the two previous columns.

Failed Interval Migrations

Represents the number of times that HSM interval migration was started, but the interval migration could not free enough space to meet a user-specified value. Excessive numbers of interval migrations and/or failed migrations indicate a potential space shortage or a threshold value in that may be incorrect.

Total Capacity Analysis

The following descriptions apply to the reports produced to analyze Total Capacity:

Primary Volume Count

The number of primary volumes found in the input data for this particular system. As previously mentioned, this value may vary, depending on the level of HSM activity for a given day.

Primary Capacity (MB)

Represents the total storage capacity of all the primary volumes included in this analysis. This value is reported in megabytes.

Percent of Primary Allocated

Represents the percentage of total primary capacity that was allocated by data sets residing on those primary volumes.

Level 1 Volume Count

The number of HSM Level 1 migration volumes found in the input data for this particular system. As previously mentioned, this value may vary, depending on the level of HSM activity for a given day.

Level 1 Capacity (MB)

Represents the total storage capacity of all the HSM Level 1 migration volumes included in this analysis. This value is reported in megabytes.

Percent of Level 1 Allocated

Represents the percentage of total Level 1 capacity that was allocated by data sets residing on those Level 1 HSM volumes.

TOTAL Volume Count

The combined number of primary volumes and HSM Level 1 migration volumes found in the input data for this particular system. As previously mentioned, this value may vary, depending on the level of HSM activity for a given day.

TOTAL Capacity (MB)

Represents the total storage capacity of all the primary and HSM Level 1 migration volumes included in this analysis. This value is reported in megabytes.

Percent of TOTAL Allocated

Represents the percentage of total primary and Level 1 capacity that was allocated by data sets residing on those volumes.

DAILY HSM DASD USAGE (STGEJC)				
CA MICS StorageMate				
ABC CORPORATION				
Excess Occupancy Analysis				
----- SYSID=IMS8 Volume Serial Number=TST068 -----				
Date	Percentage Excess Occupancy	Percentage Not Eligible	Percentage Unallocated	Total Capacity (MB)
07DEC07	27.0	42.0	31.0	1892
08DEC07	27.0	42.0	31.0	1892
09DEC07	27.0	42.0	31.0	1892
10DEC07	28.0	41.0	31.0	1892
11DEC07	14.0	55.0	31.0	1892
12DEC07	14.0	54.0	32.0	1892
13DEC07	14.0	55.0	31.0	1892
14DEC07	14.0	55.0	31.0	1892
15DEC07	13.0	55.0	32.0	1892
16DEC07	13.0	55.0	32.0	1892
17DEC07	11.0	58.0	31.0	1892
18DEC07	11.0	58.0	31.0	1892
19DEC07	11.0	58.0	31.0	1892
21DEC07	14.0	55.0	31.0	1892
22DEC07	15.0	54.0	31.0	1892
23DEC07	9.0	60.0	31.0	1892
26DEC07	12.0	57.0	31.0	1892
28DEC07	12.0	57.0	31.0	1892
29DEC07	12.0	57.0	31.0	1892
30DEC07	17.0	52.0	31.0	1892

Figure 4-25. Daily HSM DASD Usage / Summary

DASD-to-Tape Candidates (STGERA)

Source: VCADAA file at the DETAIL timespan

Function: Identifies data sets that are not stored on the most efficient storage media based on their usage patterns. For example, large DASD data sets that are used infrequently might be more cost effective if stored on tape.

Additionally, this report allows the user to specify size and usage criteria, plus the amount of data to be moved. The output will indicate those data sets that may be candidates for movement to tape.

Features: Users can limit the number of data sets appearing on the report, either in terms of the megabytes represented or the number of data sets.

Users can also create a file containing the names of all candidate data sets selected. This file can be modified by the user to create JCL and control statements to move the data sets.

SMS

Users: As part of SMS implementation planning, many users are identifying data sets that should be moved to a more efficient location, based on the size and usage pattern of the data set. This report will help perform that analysis. The name file that can be produced could be used to build JCL to move the data set.

COLUMN DESCRIPTIONS

Data Set Name

Data sets appear on this report based on their recommendation values, with the best candidates for media movement appearing first. The recommendation value is calculated as the size of the data set, multiplied by the time since the data set was last used. The best candidates for DASD-to-Tape movement are those data sets that are large in size and referenced infrequently. The number of data sets that will appear on the report is controlled by values entered on the report options screen.

Data Set Size In Megabytes

The amount of space allocated by this data set, expressed in millions of bytes (megabytes).

SYSID

The translated system identification that was derived from ORGSYSID by the CA MICS SYSID Translation process.

Volume Serial Number

The volume serial number of the storage medium.

Device Type

This value indicates the type of I/O device.

Days Since Used

The number of days since the data set has been used. This is calculated by subtracting the date of last use from the date the data was collected.

Date Last Used

The date that the data set was last referenced.

Date Created

The date that the data set was first allocated to a DASD volume.

Date Expires

The date when a data set can be deleted or overwritten by another data set.

Block Size

The maximum length of a physical block written into this data set.

Record Length

The maximum length of each logical record in this data set.

DASD-TO-TAPE CANDIDATES (STGERA)											1
CA MICS StorageMate											
ABC CORPORATION											
Data Set Name	Data Set Size In Megabytes	SYSID	Volume Serial Number	Device Type	Days Since Used	Date Last Used	Date Created	Date Expires	Block Size	Record Length	
LNR042. PRD18910. DAYS	10.7	SYS7	LNR908	3380	65	18JUN07	25JAN90	.	32760	32756	
CRLD. TST8910C. PSP. DOC. TEXT	10.1	SYS7	LNR922	3380	81	02JUN07	15MAR90	.	6160	137	
CRLD. TST8910S. DETAIL	11.9	SYS7	LNR807	3380	78	05JUN07	18MAR90	.	32760	32756	
LNR102. AN00082. SASDB	11.9	SYS7	LNR918	3380	78	05JUN07	05JUN90	.	32760	32756	
LNR117. MICSCAL. SASDB	11.9	SYS7	LNR905	3380	79	04JUN07	04JUN90	.	32760	32756	
LNR117. PAS9002. SASDB	11.9	SYS7	LNR916	3380	77	06JUN07	06JUN90	.	32760	32756	
LNR042. PRD18904. DETAIL	13.3	SYS7	LNR911	3380	69	14JUN07	14JUN89	.	32760	32756	
BBB. NDVRC1. EC1DEVEL. PRCSLIST	11.4	SYS7	LNR902	3380	85	29MAY07	16NOV88	.	6233	137	
LNR025. ISPTLIB	14.5	SYS7	LNR905	3380	67	16JUN07	19MAR90	.	6160	80	
NGALERT. TESTDATA. NETMSTR. V21. QANTEK. DATA	16.4	SYS7	LNR925	3380	63	20JUN07	20JUN90	.	23476	32756	
LNR238. ISPHLIB1	10.9	SYS7	LNR908	3380	99	15MAY07	06MAR90	.	6160	80	
LNR146. IA00085. SASDB	11.9	SYS7	LNR930	3380	91	23MAY07	23MAY90	.	32760	32756	
LNR146. APPD. SASDB	11.9	SYS7	LNR928	3380	90	24MAY07	24MAY90	.	32760	32756	
LNR042. AFT4500. MACAUTOS	14.1	SYS7	LNR929	3380	75	08JUN07	08JAN90	.	6160	80	
QCAM. ABC9004U. XYZ. DATA	15.8	SYS7	LNR923	3380	69	14JUN07	04APR90	.	9040	32760	
MASTER. DEV. VMXA. MONITOR. DATA	10.9	SYS7	LNR905	3380	104	10MAY07	10MAY90	.	23440	4092	
CONTROL. NA. MICS. RMF. DATA	18.0	SYS7	LNR804	3380	64	19JUN07	19JUN90	.	23440	32760	
CRLD. PSP8904U. DAYS	10.7	SYS7	LNR922	3380	111	03MAY07	02OCT89	.	32760	32756	
CRLD. PSP8810C. JCL. TEXT	14.9	SYS7	LNR915	3380	84	30MAY07	12NOV89	.	6160	137	
LNR042. PSP8904C. MACAUTOS	13.2	SYS7	LNR902	3380	100	14MAY07	19MAY89	.	6160	80	
LNR079. WORK. TEMP	14.2	SYS7	LNR916	3380	90	24MAY07	16MAY90	.	32760	32756	
SYS2. NCP43. SNCPMAC1	18.4	SYS7	LNR928	3380	70	13JUN07	13MAR90	.	3120	80	
LNR042. AFT4500. ISPLIB	12.5	SYS7	LNR806	3380	118	26APR07	08JAN90	.	6160	80	
CONTACT. CS4. ISPTLIB. BKUP	17.7	SYS7	LNR915	3380	84	30MAY07	27MAR90	.	6160	80	
LNR146. STATS2. SASDB	11.9	SYS7	LNR906	3380	133	11APR07	11APR90	.	32760	32756	
CRLD. PSP8904C. PSP. ISPTLIB	18.8	SYS7	LNR932	3380	84	30MAY07	06NOV89	.	6160	80	
LNR164. RECORD. TRACE. K95148	16.3	SYS7	LNR928	3380	99	15MAY07	29DEC89	.	23408	133	
LNR068. CASE. CAPDB	10.2	SYS7	LNR902	3380	160	15MAR07	15MAR90	.	32760	32756	
PROJECT. ARCDEVEL. STAGE1. DATA	10.7	SYS7	LNR910	3380	160	15MAR07	09AUG88	31DEC99	4096	0	
PROJECT. ARCDEVEL. STAGE2. DATA	10.7	SYS7	LNR910	3380	160	15MAR07	09AUG88	31DEC99	4096	0	
LNR146. ABC. SASDB	11.9	SYS7	LNR933	3380	146	29MAR07	29MAR90	.	32760	32756	
LNR146. RECD. SASDB	11.9	SYS7	LNR925	3380	148	27MAR07	27MAR90	.	32760	32756	
LNR117. APPLD. MICFTABL	17.7	SYS7	LNR933	3380	96	18MAY07	21MAR90	.	6160	80	
LNR046. SAS606. V5DETAIL	13.6	SYS7	LNR906	3380	135	09APR07	01MAR90	.	32760	32756	
CAG. TEST. DB5. SDDS. DATA	14.2	SYS7	LNR904	3380	133	11APR07	07FEB90	31DEC99	4096	0	
CAG. TEST. DB5. SDDS. DATA	14.2	SYS7	LNR911	3380	133	11APR07	07FEB90	31DEC99	4096	0	
LNR146. EXTRAD. SASDB	11.9	SYS7	LNR922	3380	160	15MAR07	15MAR90	.	32760	32756	
LNR146. STATS. SASDB	11.9	SYS7	LNR917	3380	163	12MAR07	12MAR90	.	32760	32756	
LNR046. SAS606. V6DETAIL	14.6	SYS7	LNR806	3380	134	10APR07	02APR90	.	23040	23040	
LNR146. ACCTXXX. SASDB	11.9	SYS7	LNR930	3380	169	06MAR07	06MAR90	.	32760	32756	
=====											
531.1											
N=	40										

Figure 4-26. DASD-to-Tape Candidates

DASD/HSM Data Movement / Plot (STGECE)

Source: VCADAA file at the DAYS timespan
VCA_VS file at the DAYS timespan
HSM MIG file at the DAYS timespan

Function: This facility allows the storage administrator to monitor the daily movement of data between real DASD and HSM migration volumes. Installations using HSM should experience a daily migration of less-active data to HSM level 1 and level 2 volumes. If HSM and/or SMS parameters are not set correctly, or if certain groups or users abuse HSM commands to cause data set movement, severe capacity or performance problems may result.

A summary report will analyze the first and last dates found in the input data, and create a summary of the changes between those two dates. This could be used, for example, to analyze data from a Monday and a Friday and summarize changes that occurred during the week. If potential problems are found, history reports can be used to provide more detailed information about certain groups or users.

These reports may be useful when changes have been made to HSM or SMS parameters, to validate the changes are having their intended effect on the operation of HSM.

Features: Summarization options allow you to analyze the data based on its owner (using one or more of the CA MICS accounting fields), its SMS storage class, its SMS management class, or by system (SYSID). These various levels of summarization make it easy to isolate particular problems or trends.

Data filtering allows you to include data from selected summary groups. For example, you could choose to report on specific departments groups or individuals, or on a specific SMS class.

Output reports from this facility include tabular reports, printer plots, or color graphic plots. Reporting can be based on the number of data sets, the amount of allocated space, or both.

SMS

Issues: The ability to summarize by an SMS storage or management class makes it easy to monitor the activity within each of your defined classes. The ability to filter and summarize by a specific management class, for example, could be used to verify that changes made to the HSM parameters in a management class were working as intended.

ELEMENT DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Owner

Specifies the department, group, or individual to which this chart applies. This value is specified as a report option, and represents one or more of the CA MICS accounting fields assigned to a group of data sets.

Storage Class

If summarization by storage class was requested, this indicates the name of the SMS storage class represented by this chart. All of the data sets assigned to the indicated storage class were summarized to produce this chart.

Management Class

If summarization by management class was requested, this indicates the name of the SMS management class represented by this chart. All of the data sets assigned to the indicated management class were summarized to produce this chart.

Vertical Axis (Value)

The numeric values distributed vertically down the left side of each page represent the scale of the numeric values being plotted on that page, and represent either a count of data sets or the amount of space allocated. When requesting color graphics plots, you have the option of specifying the values that will appear on this axis. This should be done with care, as all plots and groups will have to adapt to the values you specify. For printer plots, or when values are not specified, SAS will generate the values automatically, based

on the content of the plotted data. We suggest that these SAS defaults be used whenever possible, because they will avoid many potential problems.

Horizontal Axis (Date)

Plots on the horizontal axis represent the date (format ddmmyy) corresponding to the value being plotted. The number of entries appearing on the horizontal axis is determined by the number of cycles of CA MICS data specified as input to the report.

Footnotes

Three footnotes will appear at the bottom of each chart to indicate the values being plotted. The format of each will vary depending on whether you requested reporting by space or data sets. The first three listed below appear for space reporting, the last three for data set reporting.

Megabytes on Primary Storage

Indicates the amount of space, expressed in megabytes, allocated on primary DASD volumes by the data sets in this summary group.

Megabytes on HSM Level 1 Storage

Indicates the amount of space, expressed in megabytes, allocated on DFHSM migration level 1 volumes by the data sets in this summary group.

Megabytes on HSM Level 2 Storage

Indicates the amount of space, expressed in megabytes, allocated on DFHSM migration level 2 volumes by the data sets in this summary group.

Data Sets on Primary Storage

Indicates the number of data sets allocated on primary DASD volumes by the users in this summary group.

Data Sets on HSM Level 1 Storage

Indicates the number of data sets allocated on DFHSM migration level 1 volumes by the users in this summary group.

Data Sets on HSM Level 2 Storage

Indicates the number of data sets allocated on

DFHSM migration level 2 volumes by the users
in this summary group.




Figure 4-27. DASD/HSM Data Movement / Plot

DASD/HSM Data Movement / Summary (STGECE)

Source: VCADAA file at the DAYS timespan
VCA_VS file at the DAYS timespan
HSM MIG file at the DAYS timespan

Function: This facility allows the storage administrator to monitor the daily movement of data between real DASD and HSM migration volumes. Installations using HSM should experience a daily migration of less-active data to HSM level 1 and level 2 volumes. If HSM and/or SMS parameters are not set correctly, or if certain groups or users abuse HSM commands to cause data set movement, severe capacity or performance problems may result.

A summary report will analyze the first and last dates found in the input data, and create a summary of the changes between those two dates. This could be used, for example, to analyze data from a Monday and a Friday and summarize changes that occurred during the week. If potential problems are found, history reports can be used to provide more detailed information about certain groups or users.

These reports may be useful when changes have been made to HSM or SMS parameters, to validate the changes are having their intended effect on the operation of HSM.

Features: Summarization options allow you to analyze the data based on its owner (using one or more of the CA MICS accounting fields), its SMS storage class, its SMS management class, or by system (SYSID). These various levels of summarization make it easy to isolate particular problems or trends.

Data filtering allows you to include data from selected summary groups. For example, you could choose to report on specific departments groups or individuals, or on a specific SMS class.

Output reports from this facility include tabular reports, printer plots, and color graphic plots. Reporting can be based on the number of data sets, the amount of allocated space, or both.

SMS

Issues: The ability to summarize by an SMS storage or management class makes it easy to monitor the activity within each of your defined classes. The ability to filter and summarize by a specific management class, for example, could be used to verify that changes made to the HSM parameters in a management class were working as intended.

COLUMN DESCRIPTIONS**SYSID**

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Owner

Specifies the department, group, or individual to which this report applies. This value is specified as a report option, and represents one or more of the CA MICS accounting fields assigned to a group of data sets.

Storage Class

If summarization by storage class was requested, this indicates the name of the SMS storage class represented by this report. All of the data sets assigned to the indicated storage class were summarized to produce this report.

Management Class

If summarization by management class was requested, this indicates the name of the SMS management class represented by this report. All of the data sets assigned to the indicated management class were summarized to produce this report.

If reporting by data set was requested, the following report columns will be present:

Total Data Sets CHANGE

The difference in the total number of data sets allocated by this summary group between the two dates indicated. This value includes data sets found on primary storage and DFHSM migration storage. The difference is calculated using the values for the first and last dates found within the input data.

Percent Data Sets CHANGE

The percentage change in the total number of data sets allocated by this summary group between the two dates indicated. This value includes data sets found on primary storage, and DFHSM migration storage. The difference is calculated using the values for the first and last dates found within the input data.

Primary Data Sets (date)

This column appears twice--once for the first date found in the input data, and once for the last date found. The value in the column represents the total number of data sets allocated on primary storage for this summary group on the date indicated.

Primary Data Sets -CHANGE-

Represents the difference in the number of data sets reported in the previous two columns. This represents the difference in the number of data sets allocated on primary storage between the two dates shown in previous columns.

HSM Lev1 Data Sets (date)

This column appears twice--once for the first date found in the input data, and once for the last date found. The value in the column represents the total number of data sets allocated on HSM migration level 1 storage for this summary group on the date indicated.

HSM Lev1 Data Sets -CHANGE-

Represents the difference in the number of data sets reported in the previous two columns. This represents the difference in the number of data sets allocated on HSM migration level 1 storage between the two dates shown in previous columns.

HSM Lev2 Data Sets (date)

This column appears twice--once for the first date found in the input data, and once for the last date found. The value in the column represents the total number of data sets allocated on HSM migration level 2 storage for this summary group on the date indicated.

HSM Lev2 Data Sets -CHANGE-

Represents the difference in the number of data sets reported in the previous two columns. This represents the difference in the number of data sets allocated on DFHSM migration level 2 storage between the two dates shown in previous columns.

If reporting by space allocated was requested, the following report columns will be present:

Total Megabytes CHANGE

The difference in the total amount of space, in megabytes, allocated by this summary group between the two dates indicated. This value includes space allocated on primary storage, and DFHSM migration storage. The difference is calculated using the values for the first and last dates found within the input data.

Percent Megabytes CHANGE

The percentage change in the total amount of space allocated by this summary group between the two dates indicated. This value includes space allocated on primary storage and DFHSM migration storage. The difference is calculated using the values for the first and last dates found within the input data.

Primary Megabytes (date)

This column appears twice--once for the first date found in the input data, and once for the last date found. The value in the column represents the total amount of space, in megabytes, allocated on primary storage for this summary group on the date indicated.

Primary Megabytes -CHANGE-

Represents the difference in the amount of space reported in the previous two columns. This represents the difference, in megabytes, in the amount of space allocated on primary storage between the two dates shown in previous columns.

HSM Lev1 Megabytes (date)

This column appears twice--once for the first date found in the input data, and once for the last date found. The value in the column represents the amount of space, in megabytes, allocated on DFHSM migration level 1 storage for this summary group on the date indicated.

HSM Lev1 Megabytes -CHANGE-

Represents the difference, in megabytes, between the amount of allocated space reported in the previous two columns. This represents the difference in the space allocated on HSM migration level 1 storage between the two dates shown in previous columns.

HSM Lev2 Megabytes (date)

This column appears twice--once for the first date found in the input data, and once for the last date found. The value in the column represents the amount of space, in megabytes, allocated on DFHSM migration level 2 storage for this summary group on the date indicated.

HSM Lev2 Megabytes -CHANGE-

Represents the difference, in megabytes, between the amount of allocated space reported in the previous two columns. This represents the difference in the space allocated on HSM migration level 2 storage between the two dates shown in previous columns.

DASD/HSM DATA MOVEMENT (STGECE) CA MICS StorageMate - SUMMARY ABC CORPORATION											
----- SYSID=SYSA -----											
Owner	Total Data sets	Percent Data sets	Primary Data sets	Primary Data sets	Primary Data sets	HSM Lev1 Data sets	HSM Lev1 Data sets	HSM Lev1 Data sets	HSM Lev2 Data sets	HSM Lev2 Data sets	HSM Lev2 Data sets
	CHANGE	CHANGE	28MAR91	01APR91	-CHANGE-	28MAR07	01APR07	-CHANGE-	28MAR07	01APR07	-CHANGE-
TECHSUPT	7	0.6	723	724	1	1	4	3	372	375	3
SALES	-16	-0.1	16427	16380	-47	193	332	139	14029	13921	-108
FINANCE	-42	-0.4	7315	7217	-98	72	137	65	4198	4189	-9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
SYSID	-51		24465	24321	-144	266	473	207	18599	18485	-114
	=====		=====	=====	=====	=====	=====	=====	=====	=====	=====

DASD/HSM DATA MOVEMENT (STGECE) CA MICS StorageMate - SUMMARY ABC CORPORATION											
----- SYSID=SYSA -----											
Owner	Total Megabytes	Percent Megabytes	Primary Megabytes	Primary Megabytes	Primary Megabytes	HSM Lev1 Megabytes	HSM Lev1 Megabytes	HSM Lev1 Megabytes	HSM Lev2 Megabytes	HSM Lev2 Megabytes	HSM Lev2 Megabytes
	CHANGE	CHANGE	28MAR91	01APR91	-CHANGE-	28MAR07	01APR07	-CHANGE-	28MAR07	01APR07	-CHANGE-
FINANCE	107	0.2	50231	49984	-247	136	169	33	6010	6330	320
TECHSUPT	-34	-1.1	2586	2536	-50	0	2	2	504	518	14
SALES	-187	-0.3	45329	45120	-209	1601	1414	-188	11598	11807	209
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
SYSID	-115		98146	97641	-505	1737	1584	-153	18112	18655	543
	=====		=====	=====	=====	=====	=====	=====	=====	=====	=====

Figure 4-28. DASD/HSM Data Movement / Summary

DASD/HSM Data Movement / History (STGECE)

Source: VCADAA file at the DAYS timespan
VCA_VS file at the DAYS timespan
HSM MIG file at the DAYS timespan

Function: This facility allows the storage administrator to monitor the daily movement of data between real DASD and HSM migration volumes. Installations using HSM should experience a daily migration of less-active data to HSM level 1 and level 2 volumes. If HSM and/or SMS parameters are not set correctly, or if certain groups or users abuse HSM commands to cause data set movement, severe capacity or performance problems may result.

A summary report will analyze the first and last dates found in the input data, and create a summary of the changes between those two dates. This could be used, for example, to analyze data from a Monday and a Friday and summarize changes that occurred during the week. If potential problems are found, history reports may be used to provide more detailed information about certain groups or users.

These reports may be useful when changes have been made to HSM or SMS parameters, to validate the changes are having their intended effect on the operation of HSM.

Features: Summarization options allow you to analyze the data based on its owner (using one or more of the CA MICS accounting fields), its SMS storage class, its SMS management class, or by system (SYSID). These various levels of summarization make it easy to isolate particular problems or trends.

Data filtering allows you to include data from selected summary groups. For example, you could choose to report on specific departments groups or individuals, or on a specific SMS class.

Output reports from this facility include tabular reports, printer plots, or color graphic plots. Reporting can be based on the number of data sets, the amount of allocated space, or both.

SMS

Issues: The ability to summarize by an SMS storage or management class makes it easy to monitor the activity within each of your defined classes. The ability to filter and summarize by a specific management class, for example, could be used to verify that changes made to the HSM parameters in a management class were working as intended.

COLUMN DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Owner

Specifies the department, group, or individual to which this report applies. This value is specified as a report option, and represents one or more of the CA MICS accounting fields assigned to a group of data sets.

Storage Class

If summarization by storage class was requested, this indicates the name of the SMS storage class represented by this report. All of the data sets assigned to the indicated storage class were summarized to produce this report.

Management Class

If summarization by management class was requested, this indicates the name of the SMS management class represented by this report. All of the data sets assigned to the indicated management class were summarized to produce this report.

Date

Indicates the date for which the statistics on this line apply. The date is expressed in the format ddmmmyy.

TOTAL Data Sets

The total number of data sets allocated by the users of this summary group on the indicated date. This total includes data sets on primary storage, and DFHSM migration storage.

TOTAL Megabytes

The total amount of space allocated, in megabytes, by the users of this summary group on the indicated date. This total includes space allocated on primary storage and DFHSM migration storage.

Primary (Level 0) Data Sets

The total number of data sets allocated on primary storage by the users of this summary group on the indicated date.

HSM ML1 (Level 1) Data Sets

The total number of data sets allocated on DFHSM migration level 1 storage by the users of this summary group on the indicated date.

HSM ML2 (Level 2) Data Sets The total number of data sets allocated on DFHSM migration level 2 storage by the users of this summary group on the indicated date.

Primary (Level 0) Megabytes

The total amount of space allocated, in megabytes, on primary storage by the users of this summary group on the indicated date.

HSM ML1 (Level 1) Megabytes

The total amount of space allocated, in megabytes, on DFHSM migration level 1 storage by the users of this summary group on the indicated date.

HSM ML2 (Level 2) Megabytes

The total amount of space allocated, in megabytes, on HSM migration level 2 storage by the users of this summary group on the indicated date.

DASD/HSM DATA MOVEMENT (STGECE)									
CA MICS StorageMate - HISTORY									
ABC CORPORATION									
----- SYSID=SYSA Owner=FINANCE -----									
Date	*TOTAL*	*TOTAL*	Primary	HSM ML1	HSM ML2	Primary	HSM ML1	HSM ML2	
-----	Data Sets	Megabytes	(Level 0)	(Level 1)	(Level 2)	(Level 0)	(Level 1)	(Level 2)	
			Data Sets	Data Sets	Data Sets	Megabytes	Megabytes	Megabytes	
28MAR91	11585	56376	7315	72	4198	50231	136	6010	
01APR91	11543	56483	7217	137	4189	49984	169	6330	
02APR91	11439	53485	7114	122	4203	46450	157	6878	
03APR91	11481	53553	7118	123	4240	46498	160	6895	
----- SYSID=SYSA Owner=SALES -----									
Date	*TOTAL*	*TOTAL*	Primary	HSM ML1	HSM ML2	Primary	HSM ML1	HSM ML2	
-----	Data Sets	Megabytes	(Level 0)	(Level 1)	(Level 2)	(Level 0)	(Level 1)	(Level 2)	
			Data Sets	Data Sets	Data Sets	Megabytes	Megabytes	Megabytes	
28MAR91	30649	58529	16427	193	14029	45329	1601	11598	
01APR91	30633	58342	16380	332	13921	45120	1414	11807	
02APR91	30551	58030	16415	231	13905	45243	1001	11786	
03APR91	30916	58045	16884	207	13825	45410	923	11712	

Figure 4-29. DASD/HSM Data Movement / History

DASD/HSM Usage Snapshot / Chart (STGERF)

Source: VCADAA file at the DAYS timespan
VCA_VS file at the DAYS timespan
HSM MIG file at the DAYS timespan
HSBBAC file at the DAYS timespan

Function: This facility allows the storage administrator to monitor the allocation of space across both the real DASD and HSM platforms. In a typical HSM environment, as data becomes less frequently used it should migrate off of primary, or Level 0, storage and onto HSM migration level 1 storage. If the data continues to remain inactive, it should eventually migrate onto HSM level 2 storage. The HSM backup facility might also be employed to make regular incremental backups of data sets. The reports produced by this facility will give the storage administrator a concise summary of how the data sets in a particular group are allocated across these various storage levels.

These reports may be useful when investigating short-term capacity or performance problems. They can also detect problems caused by an HSM or SMS parameter not being set correctly, as well as problems caused by users not following installation standards.

Features: Summarization options allow you to analyze the data based on its owner (using one or more of the CA MICS accounting fields), its SMS storage class, its SMS management class, or its system (SYSID). These various levels of summarization make it easy to isolate particular problems or trends. For example, if you summarize by SYSID on a regular basis, you will get a general feeling of the health of your storage system. When you discover an anomaly in one of these summaries, additional reports at different summarization levels can isolate the problem to a particular group, or even a particular SMS class.

Data filtering allows you to include data from selected summary groups. For example, you could choose to report on specific departments groups or individuals, or on a specific SMS class.

Output reports from this facility include tabular

reports, printer charts, or color graphic charts. Summarization by both SMS storage class and management class may be done in the same report execution.

SMS

Issues: The ability to summarize by an SMS storage or management class makes it easy to monitor the activity within each of your defined classes. The ability to filter and summarize by a specific management class, for example, can verify that changes made to the HSM parameters in a management class are working as intended.

ELEMENT DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Owner

Specifies the department, group, or individual to which this report applies. This value is specified as a report option, and represents one or more of the CA MICS accounting fields assigned to a group of data sets.

Storage Class

If sub-grouping by storage class was requested, this indicates the name of the SMS storage class represented by this chart. All of the data sets used to create this particular chart were assigned to the indicated storage class.

Management Class

If sub-grouping by management class was requested, this indicates the name of the SMS management class represented by this chart. All of the data sets used to create this particular chart were assigned to the indicated management class.

PRIMARY (L0)

A pie slice with this label represents the amount of space allocated by data sets in this summary group on primary (level 0) storage. Numeric values may be present indicating the total amount of space allocated, expressed in megabytes (millions of bytes), and also the percentage of

all space for this user represented by this pie slice.

MIGRATE LEVEL 1

A pie slice with this label represents the amount of space allocated by data sets in this summary group on HSM migration level 1 storage. Numeric values may be present indicating the total amount of space allocated, expressed in megabytes (millions of bytes), and also the percentage of all space for this user represented by this pie slice.

MIGRATE LEVEL 2

A pie slice with this label represents the amount of space allocated by data sets in this summary group on HSM migration level 2 storage. Numeric values may be present indicating the total amount of space allocated, expressed in megabytes (millions of bytes), and also the percentage of all space for this user represented by this pie slice.

BACKUP

A pie slice with this label represents the amount of space allocated by data sets in this summary group on HSM backup storage. Numeric values may be present indicating the total amount of space allocated, expressed in megabytes (millions of bytes), and also the percentage of all space for this user represented by this pie slice.




Figure 4-30. DASD/HSM Usage Snapshot / Chart

DASD/HSM Usage Snapshot / Summary (STGERF)

Source: VCADAA file at the DAYS timespan
VCA_VS file at the DAYS timespan
HSM MIG file at the DAYS timespan
HSBBAC file at the DAYS timespan

Function: This facility allows the storage administrator to monitor the allocation of space across both the real DASD and HSM platforms. In a typical HSM environment, as data becomes less frequently used it should migrate off of primary, or Level 0, storage and onto HSM migration level 1 storage. If the data continues to remain inactive, it should eventually migrate onto HSM level 2 storage. The HSM backup facility might also be employed to make regular incremental backups of data sets. The reports produced by this facility will give the storage administrator a concise summary of how the data sets in a particular group are allocated across these various storage levels.

These reports may be useful when investigating short-term capacity or performance problems. They can also detect problems caused by an HSM or SMS parameter not being set correctly, as well as problems caused by users not following installation standards.

Features: Summarization options allow you to analyze the data based on its owner (using one or more of the CA MICS accounting fields), its SMS storage class, its SMS management class, or its system (SYSID). These various levels of summarization make it easy to isolate particular problems or trends. For example, if you summarize by SYSID on a regular basis, you will get a general feeling of the health of your storage system. When you discover an anomaly in one of these summaries, additional reports at different summarization levels can isolate the problem to a particular group, or even a particular SMS class.

Data filtering allows you to include data from selected summary groups. For example, you could choose to report on specific departments groups or individuals, or on a specific SMS class.

Output reports from this facility include tabular

reports, printer charts, or color graphic charts. Summarization by both SMS storage class and management class may be done in the same report execution.

SMS

Issues: The ability to summarize by an SMS storage or management class makes it easy to monitor the activity within each of your defined classes. The ability to filter and summarize by a specific management class, for example, could be used to verify that changes made to the HSM parameters in a management class were working as intended.

COLUMN DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Owner

Specifies the department, group, or individual to which this report applies. This value is specified as a report option, and represents one or more of the CA MICS accounting fields assigned to a group of data sets.

DFSMS Storage Class

If sub-grouping by storage class was requested, this indicates the name of the SMS storage class represented by this summary line.

DFSMS Management Class

If sub-grouping by management class was requested, this indicates the name of the SMS management class represented by this summary line.

TOTAL Data Sets

Represents the total number of data sets in this summary group allocated across all storage platforms. This includes online DASD (primary or level 0), HSM migration levels 1 and 2, and HSM backup. Note the data set count for HSM backup data sets does not include multiple backup copies of the same data set.

TOTAL Megabytes

Represents the total amount of space allocated, in

megabytes, by all the data sets in this summary group allocated across all storage platforms. This includes online DASD (primary or level 0), HSM migration levels 1 and 2, and HSM backup. Note the space allocated for HSM backup data sets includes space allocated by multiple backup copies of the same data set.

Data Sets on Levl 0 Storage

Represents the total number of data sets in this summary group that reside on primary (or level 0) volumes.

Megabytes on Levl 0 Storage

Represents the total amount of space allocated, in megabytes, by all the data sets in this summary group that reside on primary (or level 0) volumes.

Data Sets on Levl 1 Storage

Represents the total number of data sets in this summary group that reside on HSM migration level 1 storage.

Megabytes on Levl 1 Storage

Represents the total amount of space allocated, in megabytes, by all the data sets in this summary group that reside on HSM migration level 1 storage.

Data Sets on Levl 2 Storage

Represents the total number of data sets in this summary group that reside on HSM migration level 2 storage.

Megabytes on Levl 2 Storage

Represents the total amount of space allocated, in megabytes, by all the data sets in this summary group that reside on HSM migration level 2 storage.

Data Sets on Backup Storage

Represents the total number of data sets in this summary group that reside on HSM backup storage. If a data set has multiple backup copies, it is still only counted once in this total.

Megabytes on Backup Storage

Represents the total amount of space allocated, in megabytes, by all the data sets in this summary

group that reside on HSM backup storage. This amount includes space occupied by multiple copies of the same data set.

DASD/HSM USAGE SNAPSHOT (STGERF)											
CA MICS StorageMate											
ABC CORPORATION											
----- SYSID=SYSA Owner=FINANCE -----											
DFSMS Storage Class	*TOTAL* Data Sets	*TOTAL* Megabytes	Data Sets on Lev 0 Storage	Megabytes on Lev 0 Storage	Data Sets on Lev 1 Storage	Megabytes on Lev 1 Storage	Data Sets on Lev 2 Storage	Megabytes on Lev 2 Storage	Data Sets on Backup Storage	Megabytes on Backup Storage	
STANDARD	671	1644.8	298	1177.9	1	0.0	368	466.8	4	0.1	
(NONE)	428	1445.3	425	1408.0	0	0.0	3	37.3	0	0.0	
CRITICAL	1	0.0	0	0.0	0	0.0	1	0.0	0	0.0	
OWNER	1100	3090.1	723	2585.9	1	0.0	372	504.1	4	0.1	
----- SYSID=SYSA Owner=MARKETING -----											
DFSMS Storage Class	*TOTAL* Data Sets	*TOTAL* Megabytes	Data Sets on Lev 0 Storage	Megabytes on Lev 0 Storage	Data Sets on Lev 1 Storage	Megabytes on Lev 1 Storage	Data Sets on Lev 2 Storage	Megabytes on Lev 2 Storage	Data Sets on Backup Storage	Megabytes on Backup Storage	
STANDARD	7664	30732.3	3766	25368.8	69	135.6	3807	5222.7	22	5.1	
(NONE)	3155	23066.4	2973	22530.2	3	0.0	179	536.2	0	0.0	
CRITICAL	789	2583.3	576	2331.5	0	0.0	212	251.2	1	0.6	
OWNER	11608	56382.0	7315	50230.6	72	135.6	4198	6010.1	23	5.7	
----- SYSID=SYSA Owner=SALES -----											
DFSMS Storage Class	*TOTAL* Data Sets	*TOTAL* Megabytes	Data Sets on Lev 0 Storage	Megabytes on Lev 0 Storage	Data Sets on Lev 1 Storage	Megabytes on Lev 1 Storage	Data Sets on Lev 2 Storage	Megabytes on Lev 2 Storage	Data Sets on Backup Storage	Megabytes on Backup Storage	
STANDARD	23548	44601.4	10037	32102.8	193	1601.2	13174	10859.0	144	38.4	
(NONE)	7029	12084.5	6248	11420.7	0	0.0	781	663.8	0	0.0	
CRITICAL	212	1875.8	139	1803.4	0	0.0	72	72.4	1	0.0	
FASTTEST	3	5.2	2	2.3	0	0.0	1	2.8	0	0.0	
NOTCACHE	2	0.1	1	0.1	0	0.0	1	0.0	0	0.0	
OWNER	30794	58567.0	16427	45329.3	193	1601.2	14029	11598.1	145	38.4	
SYSID	43502	118039.0	24465	98145.7	266	1736.9	18599	18112.3	172	44.1	
	=====	=====	=====	=====	=====	=====	=====	=====	=====	=====	

Figure 4-31. DASD/HSM Usage Snapshot / Summary

Data Set Names (STGEFA)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMIG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan
Tape library information (external)

Function: Identifies common patterns that occur within your data set names by reporting the most frequently occurring values.

This information is useful when setting data set name standards, when calculating space usage by user, when planning to move data to another location, or when preparing to install software products that base decisions on the data set name, (such as security or space management products).

Features: Users in the process of establishing new data set naming standards can identify patterns already in use, and monitor the implementation of the new standards.

Users can more accurately report space usage by individuals and departments. Users installing automated products that are driven by data set can identify identify name patterns that allow them to start using those products sooner.

VSAM files can be included in the analysis. Multiple nodes can be combined in one execution of the report.

Report options allow analysis of up to nine nodes (1st, 2nd,...9th), a range of nodes, or the last node of the data set name.

Approximate volume counts are calculated that will assist you when planning to move groups of data sets to different types of devices.

SMS

Issues: The identification of common data set name patterns can help you classify data sets into different logical groups based on the needs of the data. Most ACS routines will also rely heavily on data set name patterns when determining the proper class name or group name to assign.

COLUMN DESCRIPTIONS

Data Set Pattern

The data set node(s) represented by this summarization. The contents of this column are determined by the values coded on the report option screen for this report, and by the values found in your data set names. The report option screen can be used to indicate which data set name node(s) are to be analyzed. A value of '** OTHERS **' as the last entry in this column represents data sets whose node(s) were not common enough to fall into any of the previous groups. The percentage of data sets that will be assigned to this group is determined by values coded on the report option screen. The occurrence of leading, trailing, or duplicate periods in this value indicates that the node(s) you specified were not present for the data set names represented by this group. For example, trying to analyze nodes 2-3 of data set name SYS1.PROCLIB would result in a value of 'PROCLIB.', while trying to use nodes 3-5 would result in a value of '..'.

Percent of all Data Sets

The percentage of all the data set names included for this analysis that had the data set name structure identified in the Data Set Pattern column.

Number of Data Sets

The number of data set names included for this analysis that had the data set name structure identified in the Data Set Pattern column.

Percent of all Space

The percentage of all the allocated space included for this analysis that had the data set name structure identified in the Data Set Pattern column.

Total Space (MB)

The total space allocated by all the data sets in this group, shown in millions of bytes (megabytes).

Number 3380-Ds Needed

The approximate number of 3380 Model D DASD

devices that would be needed to contain all of the data sets represented by this group. This value assumes that no unallocated space would be kept in reserve on any of the volumes, and that the data sets to be moved would fit exactly on the target volumes with no wasted space. For practical purposes, expand this number slightly if you are using it for actual space planning.

Number 3380-Es Needed

The approximate number of 3380 Model E DASD devices that would be needed to contain all of the data sets represented by this group. This value assumes that no unallocated space would be kept in reserve on any of the volumes, and that the data sets to be moved would fit exactly on the target volumes with no wasted space. For practical purposes, expand this number slightly if you are using it for actual space planning.

Number 3380-Ks Needed

The approximate number of 3380 Model K DASD devices that would be needed to contain all of the data sets represented by this group. This value assumes that no unallocated space would be kept in reserve on any of the volumes, and that the data sets to be moved would fit exactly on the target volumes with no wasted space. For practical purposes, expand this number slightly if you are using it for actual space planning.

Number 3390-1s Needed

The approximate number of 3390 Model 1 DASD devices that would be needed to contain all of the data sets represented by this group. This value assumes that no unallocated space would be kept in reserve on any of the volumes, and that the data sets to be moved would fit exactly on the target volumes with no wasted space. For practical purposes, expand this number slightly if you are using it for actual space planning.

Number 3390-2s Needed

The approximate number of 3390 Model 2 DASD devices that would be needed to contain all of the data sets represented by this group. This value assumes that no unallocated space would be kept in reserve on any of the volumes, and that the data sets to be moved would fit exactly on the target

volumes with no wasted space. For practical purposes, expand this number slightly if you are using it for actual space planning.

DATA SET NAMES (STGEFA)									
CA MICS StorageMate									
ABC CORPORATION									
----- System Identifier=SYSA -----									
Data Set Pattern	Percent of all Data Sets	Number of Data Sets	Percent of all Space	Total Space (MB)	Number 3380-Ds Needed	Number 3380-Es Needed	Number 3380-Ks Needed	Number 3390-1s Needed	Number 3390-2s Needed
MDEV	3.2	1176	4.1	5084	8.1	4.0	2.7	5.4	2.7
SYS1	1.8	657	3.0	3687	5.8	2.9	1.9	3.9	1.9
OSTSYS	1.4	501	1.8	2296	3.6	1.8	1.2	2.4	1.2
CUSTSAT	1.2	432	1.4	1752	2.8	1.4	0.9	1.9	0.9
MGT1RC	1.1	424	0.5	671	1.1	0.5	0.4	0.7	0.4
CSOMGT	1.1	420	0.7	832	1.3	0.7	0.4	0.9	0.4
SYS3	1.1	408	1.4	1709	2.7	1.4	0.9	1.8	0.9
OPTSYS	1.1	398	0.9	1093	1.7	0.9	0.6	1.2	0.6
DEV116	0.8	291	0.4	451	0.7	0.4	0.2	0.5	0.2
DEV046	0.8	281	0.1	179	0.3	0.1	0.1	0.2	0.1
DEV052	0.7	261	1.2	1531	2.4	1.2	0.8	1.6	0.8
DEV042	0.5	189	0.7	858	1.4	0.7	0.5	0.9	0.5
** OTHERS **	28.7	10643	17.0	21139	33.5	16.8	11.2	22.3	11.2

SYSID		16081		41282	65.5	32.8	21.8	43.6	21.8
----- System Identifier=SYSB -----									
Data Set Pattern	Percent of all Data Sets	Number of Data Sets	Percent of all Space	Total Space (MB)	Number 3380-Ds Needed	Number 3380-Es Needed	Number 3380-Ks Needed	Number 3390-1s Needed	Number 3390-2s Needed
TESTCAT	3.6	1350	6.9	8576	13.6	6.8	4.5	9.1	4.5
SYS1	3.0	1107	9.8	12155	19.3	9.6	6.4	12.8	6.4
FDRBACK	1.4	537	0.7	851	1.4	0.7	0.5	0.9	0.4
BMPTEST	1.4	533	1.5	1857	2.9	1.5	1.0	2.0	1.0
SYS5	1.1	424	1.4	1783	2.8	1.4	0.9	1.9	0.9
PRODUCT	1.1	421	4.4	5447	8.6	4.3	2.9	5.8	2.9
MSGUIDE	0.8	303	0.8	1045	1.7	0.8	0.6	1.1	0.6
ATMON	0.8	283	0.7	887	1.4	0.7	0.5	0.9	0.5
TSTMON	0.7	265	0.8	970	1.5	0.8	0.5	1.0	0.5
PERFGP	0.6	238	1.0	1206	1.9	1.0	0.6	1.3	0.6
DEVSYS	0.6	232	0.7	837	1.3	0.7	0.4	0.9	0.4
NETWRK	0.6	225	2.6	3189	5.1	2.5	1.7	3.4	1.7
MIMTEST	0.6	219	1.1	1395	2.2	1.1	0.7	1.5	0.7
DVXCM11	0.6	210	1.0	1251	2.0	1.0	0.7	1.3	0.7
CSTESTL	0.6	205	1.4	1705	2.7	1.4	0.9	1.8	0.9
** OTHERS **	38.9	14437	32.2	40050	63.5	31.8	21.2	42.3	21.2

SYSID		20989		83204	132.0	66.0	44.0	88.0	44.0
=====									
		37070		124487	197.5	98.8	65.8	131.6	65.8

Figure 4-32. Data Set Names

Data Set Performance / Chart (STGEPA)

Source: HARDVA file at the DAYS timespan
 VCADAA file at the DETAIL timespan
 VCA_VS file at the DETAIL timespan

Function: This chart estimates the average response time being received by each department or group. This is done by determining the response time being received for each active data set, based on its SYSID and VOLSER, and then finding the average of those values for each summary group. While this is not as accurate as response time analysis at the data set level, it does indicate general trends and potential problem areas. Use of this chart can help identify groups that are receiving poor response times, so that further research can be done to identify volumes or volume groups that are delivering poor service.

Features: This chart could be run regularly to monitor service times, with more detailed reporting being done when problems are detected. A wide variety of filtering, summarization and reporting options allow many views of the same data to be obtained. This allows problem research to be system-wide or to be isolated to specific users, volumes, or time periods. Data set grouping can be done using a single or range of data set nodes, a single or range of account fields, or through an external table. Both color graphics and printer graphics output are supported.

SMS

Issues: Users planning their SMS Storage Classes can approximate the average service currently being delivered to each department or application. This is important to know, because the desired response time for a group of related data sets is specified as part of the Storage Class definition. After SMS implementation is complete, this report could be used to monitor whether the expected response times are being delivered.

COLUMN DESCRIPTIONS

Date

The date on which the activity shown on this chart occurred.

Zone

If summarization by ZONE was specified on the report options screen, this identifier will list the ZONE value for which this chart applies.

Hour

If summarization by Hour was specified on the report options screen, this identifier will list the hour for which this chart applies.

Vertical Axis

The entries on the vertical axis represent the different groups of data sets for which average response times are being reported. These groups can be established using one or more of the VCA account fields, one or more of the data set name nodes, or using data extracted from a Data Set Group table. For this chart, grouping can also be requested at the DFSMS Storage Class level. When that option is selected, the value 'SC=' will precede the storage class name on the chart.

Horizontal Axis

The bar value on the horizontal axis represents the average response times being received by each reporting group, expressed in thousands of a second (milliseconds). This is calculated by dividing the sum of the response times for all the data sets in this group by the number of data sets. Response time for each data set is determined as the average response time for the volume on which the data set resides. Note that filtering criteria can be specified to limit this calculation to specific time periods.

Other

A column on the right side of the chart also reports each average response time as a numeric value.




Figure 4-33. Data Set Performance / Chart

Data Set Performance / Summary (STGEPA)

Source: HARDVA file at the DAYS timespan
VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan

Function: This report estimates the average response time being received by each department or group. This is done by determining the response time being received for each active data set, based on its SYSID and VOLSER, and then finding the average of those values for each summary group. While this is not as accurate as response time analysis at the data set level, it does indicate general trends and potential problem areas. Use of this report can help identify those groups that are receiving poor response times, and can help isolate volumes or volume groups that are delivering poor service.

Features: This report could be run regularly to monitor service times, with more detailed reporting being done when problems are detected. A wide variety of filtering, summarization and reporting options allow many views of the same data to be obtained. This allows problem research to be system-wide or to be isolated to specific users, volumes, or time periods. Data set grouping can be done using a single or range of data set nodes, a single or range of account fields, or through an external table.

Generation of the summary report is automatic.

SMS

Issues: Users planning their SMS Storage Classes can approximate the average service currently being delivered to each department or application. This is important to know, as the desired response time for a group of related data sets is specified as part of the Storage Class definition. After SMS implementation is complete, this report could be used to monitor whether the expected response times are being delivered.

COLUMN DESCRIPTIONS

Date

The date on which the activity shown on this report occurred.

Zone

If summarization by ZONE was specified on the report options screen, this identifier will list the ZONE value for which this data applies.

Hour

If summarization by Hour was specified on the report options screen, this identifier will list the hour for which this data applies.

User

Identifies the user to which this report line applies. This ownership can be established using one or more of the VCA account fields, one or more of the data set name nodes, or using data extracted from a Data Set Group table. For this report, ownership can also be requested at the DFSMS Storage Class level. When that option is selected, the value 'SC=' will precede the storage class name in this column.

Data Sets Included

Represents the number of data sets in this summarization group that were included in the analysis. To make the estimate of response time more accurate, only those data sets that have been referenced recently are included in the analysis. This criteria for determining this, along with other filtering criteria that effect this value, are specified on the report options screen.

Volumes Included

Represents the number of unique volumes on which the data sets belonging to this summarization group were found.

Average Response Time

The average response time in milliseconds being received by the data sets in this summarization group. This is calculated by dividing the sum of the response times for all the data sets in this group by the number of data sets. Response time for each data set is determined as the average response time for the volume on which the data set resides. Note that filtering criteria can be specified to limit this calculation to specific

time periods.

Best Response Time / From Volume
Reports the best response time that was being received by all the data sets in this group, and the volume that was providing that response time.

Worst Response Time / From Volume
Reports the worst response time that was being received by all the data sets in this group, and the volume that was providing that response time.

DATA SET PERFORMANCE (STGEPA)								1
CA MICS StorageMate								
ABC CORPORATION								
----- System Identifier=SYS7 Date=22AUG07 Hour=10 -----								
User	Data Sets Included	Volumes Included	Average Response Time	Best Response Time	From Volume	Worst Response Time	From Volume	
CUSTOMER-SUPPORT	24	9	20.0	14.9	WRK001	23.8	DSTLB4	
DEVELOPMENT	399	10	20.2	13.7	PROD73	31.1	WRK005	
DISTRIBUTION-LIBS	18	8	20.1	15.8	PFD003	21.7	RES704	
MARKETING-SALES	359	10	20.0	15.2	WRK002	23.0	WRK004	
SYSTEM-DATA	18	7	20.5	16.0	DXT937	19.8	TMPLIB	

Figure 4-34. Data Set Performance / Summary

Deletion/Migrati Thresholds / Chart (STGEJA)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan

Function: Graphically represents the age patterns of your data sets by that reside on primary DASD volumes.

While this information provides valuable insights into the currency of your data, it also helps you establish thresholds used by automated software to control the deletion, migration, space release, and backup of your data.

Features: Users installing DFHSM or similar automated products can choose thresholds that allow them to start using their new software sooner, without having an adverse effect on data availability.

Users already using automated products can use the charts to set new thresholds or to ensure their Existing thresholds are effective.

Report options allow the age analysis to be done for one of three values: creation date, retention period, or date of last reference. VSAM files can be included in the analysis.

Analysis can be done based on the number of data sets in each age range, or by the amount of space allocated by those data sets. Both types of reports can be produced in one report execution.

Color graphic and printer charts are supported.

SMS

Issues: Identification of common age patterns provided by the report will be useful when planning SMS management classes, because most of the services provided by the class are triggered by the age or usage of the data set. Existing SMS users can use this report to monitor their existing data and to ensure that the volumes specified in the management classes are working as expected.

ELEMENT DESCRIPTIONS

Chart Title

This analysis can be used to analyze three

different data set age values, each with its own unique meaning. The type of analysis that will be done is determined by the value entered on the report options screen, and is indicated by the last line in the title.

PLOT BASED ON DATE OF LAST REFERENCE

- Entering this value on the options screen causes the days since each data set was last referenced to be computed and reported. This is calculated by subtracting the last reference date from the date the information was collected.

PLOT BASED ON DATE OF CREATION

- Entering this value on the options screen causes the days since the data set was created to be computed and reported. This is calculated by subtracting the creation date from the date the information was collected.

PLOT BASED ON RETENTION PERIOD

- Entering this value on the options screen causes the retention days value to be computed and reported. This is calculated by subtracting the creation date of the data set from the expiration date. If no expiration date was specified, this value is set to zero.

The option chosen for the analysis is also reflected in the fourth title line of each chart. Throughout the remainder of this discussion, the word 'age' will be used when describing the operation of this report. Keep in mind that this actually can mean three different things, depending on the option chosen.

By Volume Serial Number

The volume serial number of the storage medium. Depending on the option chosen, summarization can be requested based on volume serial number. In this case, each chart will summarize the data sets found on one volume. When summarization by volume is selected, the Device Type of the volume will also appear in the chart heading.

By Data Set Prefix

This represents the first node of the data set name. Depending on the option chosen, summarization can be requested based on the data set prefix. In this case, each chart will summarize all the data sets found with the indicated prefix.

By SYSID

The translated system identification that was derived from ORGSYSID by the CA MICS SYSID Translation process. Depending on the option chosen, summarization can be requested for all the data sets residing on a single MVS system. In this case, each chart will be used to represent all the data sets found for a particular system.

By Management Class

Depending on the option chosen, summarization can be requested based on the SMS Management Class assigned to each managed data set. In this case, each chart will be used to represent all the data sets assigned to a particular SMS Management Class.

Vertical Axis

The vertical axis represents ranges of age values, starting from 1 and ending with 360 days.

Horizontal Axis

Depending on the report option chosen, each horizontal bar represents either the number of data sets in this summarization group having an age value within the range shown, or the amount of space allocated by those data sets. Because this is a cumulative chart, each horizontal bar will include the values of the previous bars as well.

Other

At the right side of the chart, several columns appear that make the chart easier to interpret. These include the data sets/space in each group, the cumulative value, the percentage of values in each group, and the cumulative percentage.

Footnotes

One of two footnotes will appear at the bottom of each chart, depending on the option selected:

NUMBER OF DATA SETS IN EACH AGE RANGE

The values appearing on the chart represent a value of 1 for each data set in each age range.

ALLOCATED SPACE (KB) IN EACH AGE RANGE

The values appearing on the chart represent the space allocated, in kilobytes, by the data sets in each age range.




Figure 4-35. Deletion/Migration/Release Thresholds / Chart

Deletion/Migrati Thresholds / Summary (STGEJA)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan

Function: Identifies the age patterns of the data sets that reside on your primary DASD volumes.

Identifies the minimum, maximum, and average dates for each group, plus a frequency range.

Features: Users installing DFSMS, DFHSM, or similar products can choose thresholds that allow them to start using their new software sooner, without having an adverse effect on data availability.

Users already using automated products can use the chart to set new thresholds or to ensure their existing thresholds are effective.

VSAM files can be included in the analysis. Panel options allow the age analysis to be done for one of three values: creation date, retention period, or date of last reference.

Analysis can be done based on the number of data sets in each age range, or by the amount of space allocated by those data sets. Both types of reports can be produced in one report execution.

SMS

Issues: Identification of common age patterns provided by the report will be useful when planning SMS management classes, because most of the services provided by the class are triggered by the age or usage of the data set. Existing SMS users can use this report to monitor their existing data and to ensure that the volumes specified in the management classes are working as expected.

COLUMN DESCRIPTIONS

Report Title

This analysis can be used to analyze three different data set age values, each with its own unique meaning. The type of analysis done is determined by the value entered on the report options screen, and is indicated by the last line in the title, as follows:

SUMMARY BASED ON DATE OF LAST REFERENCE

- Entering this value on the options screen causes the days since each data set was last referenced to be computed and reported. The value is calculated by subtracting the last reference date from the date the information was collected.

SUMMARY BASED ON DATE OF CREATION

- Entering this value on the options screen causes the days since the data set was created to be computed and reported. The value is calculated by subtracting the creation date from the date the information was collected.

SUMMARY BASED ON RETENTION PERIOD

- Entering this value on the options screen causes the retention days value to be computed and reported. The value is calculated by subtracting the creation date of the data set from the expiration date. If no expiration date was specified, this value is set to zero.

The option chosen for the analysis is also reflected in the fourth title line of each report. Throughout the remainder of this discussion, the word 'age' will be used when describing the operation of this report. Keep in mind that this actually can mean three different things, depending on the option chosen.

Volume Serial Number

The volume serial number of the storage medium. Depending on the option chosen, summarization can be requested based on volume serial number. In this case, the volume serial number will appear under this column, and the statistics in of the report line will apply to this volume. When volume summarization is selected, the Device Type of each volume will also appear in the heading.

Data Set Prefix

This represents the first node of the data set name. Depending on the option chosen,

summarization can be requested based on data set prefix. In this case, the prefix will appear under this column, and the statistics in the report line will apply to this prefix.

SYSID

The translated system identification that was derived from ORGSYSID by the CA MICS SYSID Translation process. Depending on the option chosen, summarization can be requested for all the data sets residing on a single MVS system. In this case, the SYSID of the system will appear under this column, and the statistics in the report line will apply to this system.

Management Class

Depending on the option chosen, summarization can be requested based on the SMS Management Class assigned to each managed data set. In this case, the Management Class name will appear under this column, and the statistics in the report line will apply to this class.

Valid Count

The number of data sets represented in this summarization level that had valid dates, and for which a valid age could be determined. These are the data sets that were included in the analysis for this summarization level.

Invalid Count

The number of data sets represented in this summarization level that were invalid, and for which a valid age could not be determined. These data sets are dropped, and were not included in the analysis for this summarization level.

Avg Num. Days

The average age value for all the data sets included in this level of summarization. This is calculated by dividing the sum of all the data set ages by the number of data sets in the summarization. This value is represented as a whole number of days.

Min. Num. Days

The smallest age value found within all the data sets included in this level of summarization. This value is represented as a whole number of

days.

Max. Num. Days

The largest age value found within all the data sets included in this level of summarization. This value is represented as a whole number of days.

You will see one of the following titles depending on the report options you selected:

Cumulative Percentage of Data Sets Referenced in Range of Days

This represents a cumulative percentage distribution line, based on the number of days since a data set was last referenced. For example, a value of 39 in the 181-210 column indicates that 39% of the data sets in this summary group have been referenced sometime in the past 210 days. Subtracting the previous column value from any column can be used to compute the percentage of data sets occurring in that column.

Cumulative Percentage of Space Referenced in Range of Days

This represents a cumulative percentage distribution line, based on the number of days since a data set was last referenced. For example, a value of 39 in the 181-210 column indicates that 39% of the allocated space in this summary group belongs to data sets that have been referenced sometime in the past 210 days. Subtracting the previous column value from any column can be used to compute the percentage of data sets occurring in that column.

Cumulative Percentage of Data Sets Created in Range of Days

This represents a cumulative percentage distribution line, based on the number of days since a data set was created. For example, a value of 75 in the 31-60 column indicates that 75% of the data sets in this summary group were created sometime in the past 60 days. Subtracting the previous column value from any column can be used to compute the percentage of data sets occurring in that column.

Cumulative Percentage of Space Created in Range of Days

This represents a cumulative percentage distribution line, based on the number of days since a data set was created. For example, a value of 75 in the 31-60 column indicates that 75% of the space allocated in this group is allocated to data sets created sometime in the past 60 days. Subtracting the previous column value from any column can be used to compute the percentage of data sets occurring in that column.

Cumulative Percentage of Data Sets Retained in Range of Days

This represents a cumulative percentage distribution line, based on the number of retention days specified when a data set was created. For example, a value of 58 in the 61-90 column indicates that 58% of the data sets in this summary group were given retention periods of 60 days or less when they were created. Subtracting the previous column value from any column can be used to compute the percentage of data sets occurring in that column.

Cumulative Percentage of Space Retained in Range of Days

This represents a cumulative percentage distribution line, based on the number of retention days specified when a data set was created. For example, a value of 58 in the 61-90 column indicates that 58% of the space allocated in this group belongs to data sets that were given retention periods of 60 days or less when they were created. Subtracting the previous column value from any column can be used to compute the percentage of data sets occurring in that column.

DELETION/MIGRATION/RELEASE THRESHOLDS (STGEJA)																				1
CA MICS StorageMate																				
ABC CORPORATION																				
SUMMARY BASED ON DATE OF LAST REFERENCE																				
System Identifier=SYS7 Device Type=3380																				
Volume Serial Number	Valid Count	Invalid Count	Avg. Num. Days	Min. Num. Days	Max. Num. Days	00- 07	08- 14	15- 21	22- 30	31- 60	61- 90	91- 120	121- 150	151- 180	181- 210	211- 240	241- 270	271- 300	301- 330	331- 365
LNR800	552	12	71	0	338	29	38	43	49	58	68	73	79	88	91	96	98	100	100	100
LNR801	110	7	26	0	147	50	62	68	70	83	89	96	100	100	100	100	100	100	100	100
LNR803	324	13	38	0	334	38	49	55	63	80	87	91	95	97	98	100	100	100	100	100
LNR804	345	7	36	0	331	43	53	59	67	76	87	91	95	98	99	100	100	100	100	100
LNR806	317	21	44	0	331	37	46	54	61	71	79	90	94	97	98	100	100	100	100	100
LNR807	298	12	43	0	331	37	46	51	59	73	82	89	94	98	99	100	100	100	100	100
LNR900	492	12	37	0	320	38	49	57	63	76	85	95	97	98	99	99	100	100	100	100
LNR901	549	6	75	0	316	24	32	35	40	51	58	79	81	87	93	97	99	99	100	100
LNR902	342	6	48	0	1086	38	46	52	58	70	78	94	95	97	98	99	99	99	99	99
LNR903	499	14	43	0	1084	33	44	51	58	74	83	94	95	97	99	99	100	100	100	100
LNR904	237	8	60	0	1084	35	49	52	59	68	73	79	83	92	96	98	98	98	99	99
LNR905	325	7	34	0	399	42	54	59	68	75	91	94	95	97	99	99	100	100	100	100
LNR906	195	6	41	0	1086	46	56	64	69	75	89	91	92	96	97	98	99	99	99	99
LNR907	193	4	47	0	1086	41	48	52	58	70	87	89	92	96	97	99	99	99	99	99
LNR908	319	20	41	0	481	38	48	56	61	71	84	95	96	97	98	99	99	99	99	99
LNR909	132	17	95	0	1086	33	43	48	52	61	73	83	86	88	89	93	93	93	93	94
DEVTYPE	5229	172																		
SYSID	5229	172																		

Figure 4-36. Deletion/Migration/Release Thresholds / Summary

GDG Management and Retention / By Prefix (STGEFE)

Source: VCADAA file at the DETAIL timespan
HSMIG file at the DETAIL timespan
Tape library information (external)

Function: Monitors your users' allocation and use of
Generation Data Groups (GDGs).

The report allows analysis of generation data set
allocation and usage patterns, and detects
potential problems, such a duplicate or missing
generations.

Features: User reports have three levels of summarization--
detail, base name, and data set prefix.

This report is summarized by the GDG Base Name.

The report also provides usage information, such
as the total amount of space allocated to
generation data groups at the installation.

If you are establishing or revising your standards
related to the number of generations that will be
retained, this report can help that process by
allowing you to simulate the effects of using
differing values.

This report provides a broad overview of GDG usage
by individual users or departments. This will be
useful when tracking GDG usage and allocation
patterns at a high level, and will also quickly
identify potential problems. More detailed
information about individual problems can then be
produced by rerunning the report at a lower level
of detail.

Limits: VSAM data sets cannot be included in this report,
because the GDG structure is not supported in VSAM
environments.

SMS

Issues: Users installing SMS can determine the number of
generations of a GDG that are being kept online.

This is important, because the SMS Management
class can be used to control the number of
generations that will be kept online, and the

action to take when older generations "roll off".

COLUMN DESCRIPTIONS

Prefix

Lists the high-level data set name node (prefix) that is summarized on this line of the report. All generation data groups starting with this prefix will be summarized here.

Of GDGs

The number of unique generation data groups (GDGs) that were found in the data being scanned having this prefix.

Min # Gen

The smallest number of generations found in any of the GDGs having this prefix. Low values such as 1 could indicate a potential problem, as only one copy of a particular group of data exists.

Max # Gen

The largest number of generations found in any of the GDGs having this prefix. High values indicate potential abuse of space, as many generations of similar data are being kept online.

Avg # Gen

The average number of generations being kept by all of the GDGs having this prefix. This is calculated by dividing the total number of generation data sets in this group by the total number of GDGs.

GDGs with Dups

The number of GDGs found with this prefix that contained duplicate generation numbers. For example, a value of 1 would indicate that one GDG with this prefix had one or more duplicate generations. Any value greater than zero indicates a potential problem caused by duplicate uncataloged generations.

GDGs with Gaps

The number of GDGs found with this prefix that contained gaps in the generation numbers. For example, a value of 1 would indicate that one GDG with this prefix had at least one gap in its generation numbers. Any value greater other than

zero indicates a potential problem caused by missing generations.

Min # Vol

The minimum number of volumes used by the GDGs with this prefix to store all of the generations of the same group. For example, a value of 2 indicates that all of the GDGs with this prefix have their generations spread across at least two volumes. A value such as 1 could warn of a potential data loss, because all generations of a group would be lost if that volume were lost.

Max # Vol

The maximum number of volumes used by the GDGs with this prefix to store all of the generations of the same group. For example, a value of 8 indicates that at least one of the GDGs with this prefix has generations spread across eight volumes. A high value could warn of a large number of generations being kept online, which wastes space.

Avg # Vol

The average number of volumes used by the GDGs with this prefix to store all of the generations of the same group. This is calculated by dividing the sum of the number of unique volumes used by each GDG by the number of GDGs. Low values tend to indicate potential data loss problems, while high values often indicate wasted space.

Max MB Alloc

The maximum amount of space allocated, expressed in millions of bytes (megabytes), by the GDGs with this prefix. This value reports the one GDG within this group that is using the largest amount of space.

Avg MB Alloc

The average amount of space allocated, expressed in millions of bytes (megabytes), by each of the GDGs having this prefix. This is calculated by dividing the space allocated by all of the generation data sets with this prefix by the number of GDGs.

Total MB Alloc

The total amount of space allocated, expressed in

millions of bytes (megabytes), by all of the generation data sets having this prefix.

Simulated Space Savings - Data Sets Moved

The number of data sets having this prefix that would have been rolled off the system based on the simulation values coded on the report option screen. This feature allows you to simulate the effects of keeping only a certain number of generations of each GDG online.

Simulated Space Savings - MB Saved

The amount of space, expressed in millions of bytes (megabytes) that would be released by causing certain generations having this prefix to be rolled off the system based on the simulation values coded on the report option screen. This feature allows you to simulate the effects of keeping only a certain number of generations of each GDG online.

GDG MANAGEMENT AND RETENTION (STGEFE)													11
CA MICS StorageMate - PREFIX REPORT													
ABC CORPORATION													
System Identifier=TST1													
Prefix	# Of GDGs	Min # Gen	Max # Gen	Avg # Gen	GDGs with Dups	GDGs with Gaps	Min # Vol	Max # Vol	Avg # Vol	Max MB Alloc	Avg MB Alloc	Total MB Alloc	+-+ Simulated Space Savings +-+ Data Sets Moved MB Saved
CBM00T	6	1	2	2	0	0	1	2	2	0.0	0.0	0.0	0 0.0
CIS00T	4	1	2	1	0	0	1	2	1	0.1	0.0	0.1	0 0.0
DAB10P	1	1	1	1	0	0	1	1	1	0.2	0.2	0.2	0 0.0
DCI	1	1	1	1	0	0	1	1	1	3.2	3.2	3.2	0 0.0
DCI10P	2	1	1	1	0	0	1	1	1	0.2	0.2	0.3	0 0.0
DCR6770	1	3	3	3	0	0	2	2	2	0.8	0.8	2.5	0 0.0
DMS	2	1	2	2	0	0	1	2	2	0.1	0.0	0.1	0 0.0
DPS00T	31	1	3	2	0	3	1	3	2	28.3	3.8	241.5	0 0.0
LNS00T	29	1	3	2	0	2	1	3	2	5.2	1.0	58.8	0 0.0
SMF	1	1	1	1	0	0	1	1	1	0.7	0.7	0.7	0 0.0
SMS	3	1	1	1	0	0	1	1	1	0.8	0.8	2.5	0 0.0
SSU	27	1	2	1	0	5	1	2	1	240.5	12.2	401.5	0 0.0
TESTCIF	1	1	1	1	0	0	1	1	1	0.0	0.0	0.0	0 0.0
TESTDCI	3	1	1	1	0	0	1	1	1	0.0	0.0	0.1	0 0.0
TESTDDA	37	1	2	1	2	1	1	2	1	4.8	0.2	9.8	0 0.0
TESTDKN	1	1	1	1	0	0	1	1	1	6.3	6.3	6.3	0 0.0
TESTSMS	1	1	1	1	0	0	1	1	1	138.0	138.0	138.0	0 0.0
TSTC	1	10	10	10	1	1	6	6	6	0.1	0.0	0.3	5 0.2
TSTD	2	5	6	6	1	1	3	4	4	0.0	0.0	0.1	1 0.0
-----	154				4	13						866.1	6 0.2

Figure 4-37. GDG Management and Retention / by Prefix

GDG Management and Retention / By Base Name (STGEFE)

Source: VCADAA file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
Tape library information (external)

Function: Monitors your users' allocation and use of
Generation Data Groups (GDGs).

The report allows analysis of generation data set
allocation and usage patterns, and detects
potential problems, such as duplicate or missing
generations.

Features: User reports have three levels of summarization--
detail, base name, and data set prefix.

This report is summarized by the GDG Base Name.

The report also provides usage information, such
as the total amount of space allocated to
generation data groups at the installation.

If you are establishing or revising your standards
related to the number of generations that will be
retained, this report can help that process by
allowing you to simulate the effects of using
differing values.

Limits: VSAM data sets cannot be included in this report,
as the GDG structure is not supported in VSAM
environments.

SMS

Issues: Users installing SMS can determine the number of
generations of a GDG that are being kept online.

This is important to know, because the SMS
Management class can be used to control the number
of generations that will be kept online, and the
action to take when older generations "roll off".

COLUMN DESCRIPTIONS

Base GDG Name

This lists the base Generation Data Group (GDG)
name that is summarized on this line of the
report. All generations of this group will start
with this name, suffixed with a 'GnnnnV00' value,

where 'nnnn' represents the generation number.

of GENS

The number of generations of this GDG that were found in the data being scanned. For example, if the following data sets were found in the scanned data:

```
PAYROLL.BACKUP.G0045V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0048V00
```

the GDG Base Name would be PAYROLL.BACKUP, and the value of this column would be 4, indicating four data sets were found.

Min GEN

The minimum (oldest) generation number found for this GDG in the data being scanned. For example, if the following data sets were found in the scanned data:

```
PAYROLL.BACKUP.G0045V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0048V00
```

the GDG Base Name would be PAYROLL.BACKUP, and the value of this column would be 45, indicating that generation 45 was the oldest generation found.

Max GEN

The maximum (newest) generation number found for this GDG in the data being scanned. For example, if the following data sets were found in the scanned data:

```
PAYROLL.BACKUP.G0045V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0048V00
```

the GDG Base Name would be PAYROLL.BACKUP, and the value of this column would be 48, indicating that generation 48 was the most recent generation found.

Dupe GENS

The number of duplicate generations found for this GDG in the data being scanned. For example, if the following data sets were found in the scanned data:

```
PAYROLL.BACKUP.G0045V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0048V00
```

the GDG Base Name would be PAYROLL.BACKUP, and the value of this column would be 1, because generation 46 is duplicated in the data. Any value other than zero should be considered a potential error. But first make sure you do not have duplicate data being read by the report.

of Gaps

The number of gaps found in the generation numbers for this GDG in the data being scanned. For example, if the following data sets were found in the scanned data:

```
PAYROLL.BACKUP.G0045V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0046V00  
PAYROLL.BACKUP.G0048V00
```

the GDG Base Name would be PAYROLL.BACKUP, and the value of this column would be 1, because generation 47 is missing from the data. Any value other than zero should be considered a potential error.

of Vols

The number of unique volumes on which the generations of this GDG reside. This number is an indication of how well the different members of a GDG are distributed across your volumes. Take note of GDGs that only reside on one volume, as that could cause a potential loss of all data if that volume were lost.

Min MB Alloc

The minimum amount of space allocated, expressed in millions of bytes (megabytes), by the smallest generation within this GDG.

Max MB Alloc

The maximum amount of space allocated, expressed in millions of bytes (megabytes), by the largest generation within this GDG.

Avg MB Alloc

The average amount of space allocated, expressed in millions of bytes (megabytes), by each of the generations within this GDG. This is calculated by dividing the total space allocated by all generations by the number of generations.

Total MB Alloc

The total amount of space allocated, expressed in millions of bytes (megabytes), by all of the generations within this GDG.

Simulated Space Savings - Data Sets Moved

The number of generations within this GDG that would have been rolled off the system based on the simulation values coded on the report option screen. This feature allows you to simulate the effects of keeping only a certain number of generations of each GDG online.

Simulated Space Savings - MB Saved

The amount of space, expressed in millions of bytes (megabytes) that would be released by causing certain generations within this GDG to be rolled off the system based on the simulation values coded on the report option screen. This feature allows you to simulate the effects of keeping only a certain number of generations of each GDG online.

GDG MANAGEMENT AND RETENTION (STGEFE)													7
CA MICS StorageMate - BASE NAME REPORT													
ABC CORPORATION													
----- System Identifier=TS1 -----													
Base GDG Name	# of GENS	Min GEN #	Max GEN #	Dupe GENS	# of Gaps	# of Vols	Min MB Alloc	Max MB Alloc	Avg MB Alloc	Total MB Alloc	+-Simulated Data Sets Moved	Space Savings- MB Saved	
CBM00T.CARDTAPE.NCB00110.MS000IN0	2	1	2	0	0	2	0.0	0.0	0.0	0	0	0.0	
CBM00T.CARDTAP2.NCB00110.MS000IN0	1	2	2	0	0	1	0.0	0.0	0.0	0	0	0.0	
CBM00T.CBFCOUT.NCB00110.MS000IN0	1	2	2	0	0	1	0.0	0.0	0.0	0	0	0.0	
CBM00T.CBPINEXT.NCB00110.MS000IN0	2	1	2	0	0	2	0.0	0.0	0.0	0	0	0.0	
CBM00T.CBPUNLD.NCB00110.MS000IN0	2	1	2	0	0	2	0.0	0.0	0.0	0	0	0.0	
CBM00T.MLERTAPE.NCB00110.MS000IN0	2	1	2	0	0	1	0.0	0.0	0.0	0	0	0.0	
CIS00T.BCPRNT00.RPT#0001	1	2	2	0	0	1	0.0	0.0	0.0	0	0	0.0	
CIS00T.BCPRNT00.RPT#0002	1	1	1	0	0	1	0.0	0.0	0.0	0	0	0.0	
CIS00T.DLDARPGS.ICI00211.MS000IN0	1	3	3	0	0	1	0.1	0.1	0.1	0	0	0.0	
CIS00T.DLRPTA00.RPT#0001	2	1	2	0	0	2	0.0	0.0	0.0	0	0	0.0	
DAB10P.DCIPTSTB.RPT#B001	1	1	1	0	0	1	0.2	0.2	0.2	0	0	0.0	
DCI.PRMST135.DCIM050F.FNB1GDG	1	183	183	0	0	1	3.2	3.2	3.2	3	0	0.0	
DPS00T.DLCRBUGS.IDP00621.MS000IN0	1	1	1	0	0	1	0.0	0.0	0.0	0	0	0.0	
DPS00T.DLRPT000.RPT#0006	3	2	4	0	0	3	0.0	0.0	0.0	0	0	0.0	
DPS00T.DLRPT000.RPT#0008	1	4	4	0	0	1	0.0	0.0	0.0	0	0	0.0	
DPS00T.DPAAXTGS.IDP00621.MS000IN0	1	2	2	0	0	1	0.0	0.0	0.0	0	0	0.0	
DPS00T.DPBULKGS.IDP00621.MS000IN0	3	4	6	0	0	3	11.4	11.4	11.4	34	0	0.0	
DPS00T.DPCOFRGS.IDP00621.MS000IN0	3	4	6	0	0	3	0.7	0.7	0.7	2	0	0.0	
DPS00T.DPCONRGS.IDP00621.MS000IN0	3	1	5	0	1	2	0.2	28.3	18.9	57	0	0.0	
DPS00T.DPCRIMGS.IDP00621.MS000IN0	3	4	6	0	0	2	9.2	9.2	9.2	28	0	0.0	
DPS00T.DPDARPGS.IDP00621.MS000IN0	1	3	3	0	0	1	2.2	2.2	2.2	2	0	0.0	
DPS00T.DPDLBLS.IDP00621.MS000IN0	2	4	5	0	0	2	2.8	2.8	2.8	6	0	0.0	
DPS00T.DPEXPSGS.IDP00621.MS000IN0	2	4	5	0	0	2	11.4	11.5	11.5	23	0	0.0	
DPS00T.DPTRCVCH.NDP00621.MS000IN0	1	9	9	0	0	1	0.0	0.0	0.0	0	0	0.0	
DPS00T.DPTRRTV0.RPT#0006	1	4	4	0	0	1	0.0	0.0	0.0	0	0	0.0	
DPS00T.FLOATINP.NDP00621.MS000IN0	2	5	6	0	0	2	0.0	0.0	0.0	0	0	0.0	
DPS00T.MEMOTRAN.NDP00621.MS000IN0	3	2	4	0	0	3	0.0	0.0	0.0	0	0	0.0	
DPS00T.MLTPOFLD.NDP00621.MS000IN0	1	1	1	0	0	1	0.0	0.0	0.0	0	0	0.0	
DPS00T.OFFLDINP.NDP00621.MS000IN0	3	5	7	0	0	3	0.1	0.1	0.0	0	0	0.0	
DPS00T.PENDUNLD.NDP00621.MS000IN0	3	1	8	0	1	2	0.0	0.0	0.0	0	0	0.0	
DPS00T.PEPINPUT.NDP00621.MS000IN0	2	5	6	0	0	2	0.0	0.0	0.0	0	0	0.0	
DPS00T.RETTAPE0.NDP00621.MS000IN0	3	6	8	0	0	3	0.0	0.0	0.0	0	0	0.0	
DPS00T.RIMINPUT.NDP00621.MS000IN0	3	2	6	0	1	2	0.0	0.0	0.0	0	0	0.0	
LNS00T.CILNACPG.ILN00330.MS000IN0	2	3	4	0	0	2	3.7	3.7	3.7	7	0	0.0	
LNS00T.COKEYWRD.NLN00330.MS000IN0	3	4	6	0	0	2	1.5	1.5	1.5	5	0	0.0	
LNS00T.COUPON00.ILN00330.MS000IN0	3	2	4	0	0	3	5.2	5.2	5.2	15	0	0.0	
SYSID	253				5	14				866	6	0.2	

Figure 4-38. GDG Management and Retention / Base Name

GDG Management and Retention / Detail (STGEFE)

Source: VCADAA file at the DETAIL timespan
HSMMIG file at the DETAIL timespan
Tape library information (external)

Function: Monitors your users' allocation and use of
Generation Data Groups (GDGs) at a detailed level.

The report allows analysis of generation data set
allocation and usage patterns, and detects
potential problems, such as duplicate or missing
generations.

Features: User reports have three levels of summarization--
detail, base name, and data set prefix.

This report is summarized by the GDG Base Name.

The report also provides usage information, such
as the total amount of space allocated to
generation data groups at the installation.

If you are establishing or revising your standards
related to the number of generations that will be
retained, this report can help that process by
allowing you to simulate the effects of using
differing values.

When you run this report using prefix or base name
summarizations, you will often detect potential
problems at the prefix or GDG level. This
detailed reporting option can then be used to
provide further insight into those potential
problems.

Limits: VSAM data sets cannot be included in this report,
as the GDG structure is not supported in VSAM
environments.

SMS

Issues: Users installing SMS can determine the number of
generations of a GDG that are being kept online.

This is important to know, because the SMS
Management class can be used to control the number
of generations that will be kept online, and the
action to take when older generations "roll off".

COLUMN DESCRIPTIONS

Data Set Name

The name of the data set.

Volume Serial Number

The volume serial number of the storage medium.
For data sets that reside on HSM storage, this will have a value of 'HSM LVx' where 'x' will be 1 or 2 representing the migration level.

Roll Off Flag

An asterisk (*) in this column indicates this data set would have been a candidate for roll-off based on the simulation values you coded on the report options screen. This simulation feature allows you to specify the number of generations that should be kept online for each Generation Data Group. Generations which exceed this number are termed 'rolled-off', and simulates the effect of migrating, moving, or deleting the data set.

Data Set Org

The organization of the data.

Extents

The number of non-contiguous segments of space allocated to the data set on one volume. This value will be blank for data sets residing on tape or HSM migration volumes.

Megabytes Allocated

The amount of space allocated by this data set, in megabytes.

Megabytes Used

The amount of space actually used by this data set, in megabytes. For data sets residing on tape or HSM, this will be the same as megabytes allocated (unless you have specified the PRIMARY space reporting option for HSM data sets).

Create Date

The date that the data set was first allocated to a DASD volume.

Date Last Used

The date that the data set was last referenced.

GDG MANAGEMENT AND RETENTION (STGEFE)								1
CA MICS StorageMate - DETAIL								
ABC CORPORATION								
----- System Identifier=TST1 -----								
Data Set Name	Volume Serial Number	Roll Off Flag	Data Set Org	Extents	Megabytes Allocated	Megabytes Used	Create Date	Date Last Used
CBM00T.CARDTAPE.NCB00110.MS000IN0.G0001V00	VH1001		PS	0	0.00	0.00	16FEB90	16FEB07
CBM00T.CARDTAPE.NCB00110.MS000IN0.G0002V00	VH1003		PS	0	0.00	0.00	19MAR90	19MAR07
CIS00T.DLDARPGS.ICI00211.MS000IN0.G0003V00	VH1002		PS	1	0.05	0.05	11JAN90	12JUL07
CIS00T.DLRPTA00.RPT#0001.G0001V00	V10001		PS	0	0.00	0.00	13DEC89	13OCT07
CIS00T.DLRPTA00.RPT#0001.G0002V00	VH1001		PS	0	0.00	0.00	11JAN90	11JAN07
DAB10P.DCIPTSTB.RPT#B001.G0001V00	VS1002	***		1	0.15	0.00	12MAR90	.
DCI.PRMST135.DCIM050F.FNB1GDG.G0183V00	VH1001		PS	1	2.08	2.08	31JUL90	31JUL07
DCI10P.DCIPTSTA.RPT#0002.G0001V00	VS1004	***		1	0.15	0.00	07FEB90	.
DCI10P.DCIPTSTB.RPT#B001.G0001V00	VH1001	***		1	0.15	0.00	09MAR90	.
DCR6770.TEST.GDG.G0001V00	VS1002	***		1	0.75	0.00	18MAY90	.
DCR6770.TEST.GDG.G0002V00	VS1004	***		1	0.75	0.00	18MAY90	.
DCR6770.TEST.GDG.G0003V00	VS1004	***		1	0.75	0.00	18MAY90	.
DMS.DMS0P003.DMSI050I.MS000IN0.G0049V00	VH1002		PS	1	0.05	0.05	07MAR90	09MAR07
DMS.TESTDUMP.DMSI050T.FNB1DATA.G0003V00	VH1004		PS	0	0.00	0.00	07DEC88	09JAN07
DMS.TESTDUMP.DMSI050T.FNB1DATA.G0004V00	V10001		PS	1	0.05	0.05	07DEC88	09JAN07
DPS00T.CPCINPT.NDP00621.MS000IN0.G0005V00	V10001		PS	0	0.00	0.00	29JAN90	29JAN07
DPS00T.CPCINPT.NDP00621.MS000IN0.G0006V00	VH1003		PS	0	0.00	0.00	30JAN90	30JAN07
DPS00T.CPCINPT.NDP00621.MS000IN0.G0007V00	V10004		PS	1	0.30	0.30	12JUL90	13JUL07
DPS00T.DATACHG0.NDP00621.MS000IN0.G0005V00	V10001		PS	0	0.00	0.00	29JAN90	29JAN07
DPS00T.DATACHG0.NDP00621.MS000IN0.G0006V00	VH1004		PS	0	0.00	0.00	30JAN90	30JAN07
DPS00T.DLCRBUGS.IDP00621.MS000IN0.G0001V00	V10001		PS	1	0.05	0.05	28DEC89	28OCT07
DPS00T.DLRPT000.RPT#0006.G0002V00	V10003		PS	0	0.00	0.00	28JAN90	28JAN07
DPS00T.DLRPT000.RPT#0006.G0003V00	V10004		PS	0	0.00	0.00	28JAN90	28JAN07
DPS00T.DLRPT000.RPT#0006.G0004V00	VH1002		PS	0	0.00	0.00	30JAN90	30JAN07
DPS00T.DLRPT000.RPT#0008.G0004V00	VH1001		PS	0	0.00	0.00	30JAN90	30JAN07
DPS00T.DPAAXTGS.IDP00621.MS000IN0.G0002V00	V10001		PS	1	0.05	0.05	28DEC89	28OCT07
DPS00T.DPBULKGS.IDP00621.MS000IN0.G0004V00	VH1004	***		1	10.05	0.00	29JAN90	.
DPS00T.DPBULKGS.IDP00621.MS000IN0.G0005V00	V10001	***		1	12.05	0.00	29JAN90	.
DPS00T.DPBULKGS.IDP00621.MS000IN0.G0006V00	V10003	***		1	12.05	0.00	30JAN90	.
DPS00T.DPCOFRGS.IDP00621.MS000IN0.G0004V00	VH1004	***		1	0.65	0.00	29JAN90	.
DPS00T.DPCOFRGS.IDP00621.MS000IN0.G0005V00	V10001	***		1	0.75	0.00	29JAN90	.
DPS00T.DPCOFRGS.IDP00621.MS000IN0.G0006V00	V10003	***		1	0.75	0.00	30JAN90	.
DPS00T.DPCONRGS.IDP00621.MS000IN0.G0001V00	V10001		PS	1	0.25	0.25	28DEC89	28OCT07
DPS00T.DPCONRGS.IDP00621.MS000IN0.G0004V00	V10001	***		1	29.85	0.00	29JAN90	.
DPS00T.DPCONRGS.IDP00621.MS000IN0.G0005V00	VH1002	***		1	24.85	0.00	29JAN90	.
DPS00T.DPCRIMGS.IDP00621.MS000IN0.G0004V00	V10003	***		1	9.70	0.00	29JAN90	.
DPS00T.DPCRIMGS.IDP00621.MS000IN0.G0005V00	V10001	***		1	9.70	0.00	29JAN90	.
DPS00T.DPCRIMGS.IDP00621.MS000IN0.G0006V00	V10003	***		1	9.70	0.00	30JAN90	.
DPS00T.DPDARPGS.IDP00621.MS000IN0.G0003V00	VH1001		PS	1	1.95	1.95	28JAN90	30JAN07
DPS00T.DPDLBLS.IDP00621.MS000IN0.G0004V00	V10003	***		1	2.95	0.00	29JAN90	.

Figure 4-39. GDG Management and Retention / Detail

I/O Activity by Workload / Chart (STGEPC)

Source: BATWDA file at the DAYS timespan
HARDVA file at the DAYS timespan
WLMSEC file at the DAYS timespan

Function: This chart provides a graphic view of the I/O activity against each DASD volume, based on the workload EXCP counts stored by CA MICS in the BATWDA file.

Each user of the volume is identified by the WLM Service Class name, which is then displayed on the chart.

Users planning for storage products that support volume pooling can use this report to identify usage patterns that must be known when defining volume groups if proper performance is to be received for each group.

Features: Three different CA MICS files are combined to produce this chart. The user can set the summarization level to day, zone, or hour, so that both broad and narrow views of the data can be obtained. This chart is useful in detecting current or potential performance problems at the volume level, caused by an improper mix of users on the volume.

Volumes with low levels of I/O activity can, optionally, be excluded from the analysis.

SMS

Issues: Identifying the primary users of each volume during critical time periods is an important step when planning SMS storage groups. This report will enable you to group volumes with similar workloads together, thus avoiding potential performance problems caused by conflicting users.

COLUMN DESCRIPTIONS

Date

The date on which the activity shown on this chart occurred.

Volume

The volume serial number of the storage medium.

Zone

If summarization by ZONE was specified on the report options screen, the data on this chart will only apply for this ZONE.

Hour

If summarization by Hour was specified on the report options screen, the data on this chart will only apply for this hour.

Pie Slices

Each slice represents the number of I/O requests for this volume received during this measurement interval that originated from a particular workload. The percentage value represents the percentage of all requests received that originated from each workload.

Pie Labels

Each slice identifies the name of the workload that issued I/O requests to this volume. Similar tasks in MVS are assigned to a common WLM Service Class, which is controlled and reported by means of a service class name. A value of '(SYSTEM)' (?) in the label is used to represent uncaptured I/O requests. This refers to I/O requests that are captured in the totals from RMF for the device, but are not captured by SMF and assigned to the proper user. This is a 'normal' occurrence for all MVS systems, as some work performed by the operating system is considered to be overhead, and is not captured in job accounting.




Figure 4-40. I/O Activity by Workload / Chart

I/O Activity by Workload / Detail (STGEPC)

Source: BATWDA file at the DAYS timespan
HARDVA file at the DAYS timespan
WLMSEC file at the DAYS timespan

Function: Monitors the I/O activity of each DASD volume, based on the workload EXCP counts stored by CA MICS in the BATWDA file.

Each user of the volume is identified by the WLM Service Class name. This gives the storage administrator the ability to monitor the work being done on each volume during various time periods.

Users planning for storage products that support volume pooling can use this report to identify usage patterns that must be known when defining volume groups if proper performance is to be received for each group.

Features: Detects current or potential performance problems at the volume level, caused by an improper mix of users on the volume.

Three different CA MICS files are combined to produce this report. The user can set the summarization level to day, zone, or hour, so that both broad and narrow views of the data can be obtained.

Volumes with low levels of I/O activity can, optionally, be excluded from the analysis.

SMS

Issues: Identifying the primary users of each volume during critical time periods is an important step when planning SMS storage groups. This report will enable you to group volumes with similar workloads together, thus avoiding potential performance problems caused by conflicting users.

COLUMN DESCRIPTIONS

Date

The date on which the activity shown on this report occurred.

Volume Serial Number

The volume serial number of the storage medium.

Zone

If summarization by ZONE was specified on the report options screen, this column will list the ZONE value for which this data applies.

Hour

If summarization by Hour was specified on the report options screen, this column will list the hour for which this data applies.

Service Class Name

The name of the WLM Service Class that issued I/O requests against this volume during this measurement period. Similar tasks in MVS are assigned to a common WLM Service Class, which is controlled and reported by means of a service class name.

Report Class Name

The name of the WLM Report Class that issued I/O requests against this volume during this measurement period. Similar tasks in MVS are often grouped into a optional WLM Report Class, which is used for reporting service and resource consumption. Each WLM Report Class is known by a unique name.

I/Os Per Sec.

The average number of I/O requests per second that were directed to this volume from this workload. This is calculated by dividing the number of I/O requests by the length of the RMF measurement interval, in seconds. This column is sorted so that the workloads issuing the most requests will appear first.

Percent of Total I/Os

The percentage of all the I/O requests for this volume received during this measurement interval that originated from this workload. This column is sorted so that the workloads issuing the most requests will appear first. For the '(SYSTEM)' workload, this value represents the percentage of the I/O for this volume that is not being captured by MVS.

Number of I/O Requests

The total number of I/O requests for this volume received during this measurement interval that originated from this workload. This column is sorted so that the workloads issuing the most requests will appear first.

Device Addr.

The physical address of an I/O device.

Device Type

The type of I/O device.

Storage Group

In an installation running under Data Facility Product (MVS/DFP) Version 3 with the Data Facility Storage Management Subsystem (DFSMS) active, STORNAME contains the Storage Group Name that is the equivalent of a DASD pool name. It represents a group of physical DASD volumes to be managed by DFSMS and is defined by the installation's DASD Storage Administrator.

I/O ACTIVITY BY WORKLOAD (STGEPC)										1					
CA MICS StorageMate - DETAIL															
ABC CORPORATION															
----- System Identifier=SYS7 Date=22AUG90 -----															
Volume Serial Number	Hour	Service Class Name	Report Class Name	I/Os Per Sec.	Percent of Total I/Os	Number of I/O Requests	Device Addr.	Device Type	Storage Group						
LNR907	10	TSO/DEF	NONSYSUS	0.423	39.0	1522	14D	3380							
		(SYSTEM)		0.313	28.8	1125									
		BAT/DEF		0.217	20.0	782									
		STCLONG		0.130	12.0	467									
		STCLONG		0.001	0.1	3									
		STCSHORT		0.001	0.1	3									
				-----	-----	-----									
				1.084	100.0	3903									
LNR907	11	TSO/DEF	NONSYSUS	0.924	70.9	3325	14D	3380							
		BAT/DEF		0.154	11.8	555									
		(SYSTEM)		0.154	11.8	553									
		STCLONG		0.070	5.4	251									
		STCLONG		0.001	0.1	3									
						-----				-----	-----				
				1.302	100.0	4687									
LNR908	10	BAT/DEF	NONSYSUS	0.640	35.7	2303	14F	3380							
		STC/LOW		0.431	24.0	1552									
		(SYSTEM)		0.396	22.1	1426									
		TSO/DEF		0.326	18.2	1172									
		STCLONG		0.001	0.0	2									
						-----				-----	-----				
				1.793	100.0	6456									
LNR908	11	TSO/DEF	NONSYSUS	0.180	35.8	647	14F	3380							
		BAT/DEF		0.168	33.3	603									
		(SYSTEM)		0.153	30.5	551									
		STCLONG		0.002	0.4	8									
						-----				-----	-----				
						0.503				100.0	1809				

Figure 4-41. I/O Activity by Workload / Detail

Monthly Application Resource Usage / Plot (STGERE)

Source: VCADAA file at the MONTHS timespan
VCA_VS file at the MONTHS timespan
HSM MIG file at the MONTHS timespan
HSBBAC file at the MONTHS timespan

Depending on your CA MICS installation options, these files may not be active in the MONTHS timespan. Check with your CA MICS Administrator if you have problems when trying to use this report.

Function: This facility allows the Storage Administrator to monitor key storage indicators on a month-by-month basis. These indicators enable accurate projections of future storage needs, both at the application level and the system level. The indicators reported by this facility are also useful for detecting trends that could indicate potential problems or areas in which storage resources are not being used efficiently. These reports, along with the Daily Application Resource Usage reports, provide both a long view and a short view of storage capacity and performance problems.

The ability to process data from multiple sources makes the reports produced even more valuable. Options allow the user to target a single storage medium, such as real DASD or HSM backup storage, or to group data from multiple media on the same chart. This second type of reporting is useful when analyzing the flow of data across multiple types of storage media.

The types of data that can be selected for reporting include the number of data sets, amount of space being allocated, percentages of data sets and space in certain categories, and the amount of allocated versus used space.

Features: Summarization options allow you to analyze the data based on its owner (using one or more of the CA MICS accounting fields), its SMS storage class, its SMS management class, or its location (such as real DASD, or HSM level 1 migration).

Data filtering allows you to include data from

selected summary groups. For example, you could choose to report on specific departments, groups or individuals, or on a specific SMS class.

Output reports from this facility include tabular reports, printer plots, or color graphic plots. When plots are requested, up to nine different elements may be plotted on the same chart.

SMS

Issues: The ability to summarize by an SMS storage or management class makes it easy to monitor the long-term activity within each of your defined classes. Many of the elements you can select for reporting are useful for monitoring an SMS migration effort, or to make sure SMS is operating correctly after migration is complete.

ELEMENT DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Media

Storage Class

Management Class

Group

The top line of each page also indicates the summary group to which the values on this plot apply. Depending on the report option chosen, this summarization will be done by media (DASD, VSAM, MIGRATE, or BACKUP), SMS storage class, SMS management class, or based on the CA MICS accounting codes in the data. The label to the left of the summary value will be one of the values shown above, depending on the summarization option chosen.

Vertical Axis (Value)

The numeric values distributed vertically down the left side of each page represent the scale of the numeric values being plotted on that page. When requesting color graphics plots, you have the option of specifying the values that will appear on this axis. This should be done with care, because all plots and groups will have to adapt to the values you specify. For printer plots, or

when values are not specified, SAS will generate the values automatically, based on the content of the plotted data. We suggest that these SAS defaults be used whenever possible, because they will avoid many potential problems. Care should also be taken when specifying multiple elements on one plot, so that the elements chosen will have a similar range of values.

Horizontal Axis (yy/mm)

Plots on the horizontal axis represent the year and month corresponding to the value being plotted.

The number of entries appearing on the horizontal axis is determined by the number of cycles of CA MICS data specified as input to the inquiry.

Footnotes

The footnote(s) section of each graph is important, because it indicates the names of the data element(s) that are plotted on the graph. Up to 9 elements may be plotted on the same graph, in which case 9 footnotes will appear at the bottom of each page. For printer plots, an alphabetic character A-I will appear to the left of each description, identifying the character used on the plot to represent that particular element. On color graphic plots, the letter does not appear, because the color of the footnote should correspond with the color of the plot point/line on the graph. If different colors do not appear for each point and footnote, your MICS options should be modified to specify different colors. A brief description of the possible values in the footnote section, and what those values represent, is given below:

Data Set Count: Total

Represents the total number of data sets found in the input data for this particular summary group.

Data Set Count: SMS-Managed

Represents the number of data sets found in the input data for this particular summary group that are managed by SMS.

Data Set Count: Unmanaged

Represents the number of data sets found in the input data for this particular summary

group that are not managed by SMS.

Percentage of Data Sets: Managed

Represents the percentage of all the data sets found in the input data for this particular summary group that are managed by SMS.

Percentage of Data Sets: Unmanaged

Represents the percentage of all the data sets found in the input data for this particular summary group that are not managed by SMS.

Space Allocated (MB): Total

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group.

Space Allocated (MB): Managed

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are managed by SMS.

Space Allocated (MB): Unmanaged

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Space: Managed

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are managed by SMS.

Percentage of Space: Unmanaged

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not managed by SMS.

Data Set Count: Non-VSAM

Represents the number of data sets found in the input data for this particular summary group that are not VSAM data sets.

Data Set Count: VSAM

Represents the number of data sets found in the input data for this particular summary

group that are VSAM data sets.

Percentage of Data Sets: Non-VSAM

Represents the percentage of data sets found in the input data for this particular summary group that are not VSAM data sets.

Percentage of Data Sets: VSAM

Represents the percentage of data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Space Used: All Data

Represents the percentage of space allocated versus used for all the data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the data sets in this group by the total tracks allocated. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Managed

Represents the percentage of space allocated versus used for the SMS-managed data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the managed data sets in this group by the total tracks allocated by the managed data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Unmanaged

Represents the percentage of space allocated versus used for the unmanaged data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the unmanaged data sets in this group by the total tracks allocated by the unmanaged data sets. Note that HSM data sets are not included in this calculation.

Space Allocated (MB): Non-VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not VSAM data sets.

Space Allocated (MB): VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Space: Non-VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not VSAM data sets.

Percentage of Space: VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are VSAM data sets.

Percentage of Space Used: Non-VSAM

Represents the percentage of space allocated versus used for the non-VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the non-VSAM data sets in this group by the total tracks allocated by the non-VSAM data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: VSAM

Represents the percentage of space allocated versus used for the VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the VSAM data sets in this group by the total tracks allocated by the VSAM data sets. Note that HSM data sets are not included in this calculation.

Data Set Count: HSM Level 1

Represents the number of data sets found in the input data for this particular summary group that reside on HSM migration level 1 storage.

Data Set Count: HSM Level 2

Represents the number of data sets found in the input data for this particular summary group that reside on HSM migration level 2 storage.

Data Set Count: HSM Backup

Represents the number of data sets found in the input data for this particular summary group that reside on HSM backup storage. Note that multiple backup copies of the same data set do not effect this count, nor is there any attempt made to indicate which of these data sets may also reside on primary storage.

Space Allocated (MB): HSM Level 1

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on HSM migration level 1 storage.

Space Allocated (MB): HSM Level 2

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on HSM migration level 2 storage.

Space Allocated (MB): HSM Backup

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on HSM backup storage. This total includes space used for multiple backup copies of the same data set.

Percentage of Data Sets: HSM Level 1

Represents the percentage of data sets found in the input data for this particular summary group that reside on HSM migration level 1 storage.

Percentage of Data Sets: HSM Level 2

Represents the percentage of data sets found in the input data for this particular summary group that reside on HSM migration level 2 storage.

Percentage of Data Sets: HSM Backup

Represents the percentage of data sets found in the input data for this particular summary group that reside on HSM backup storage. Note that multiple backup copies of the same data set do not effect this count, nor is there any attempt made to indicate which of these data sets may also reside on primary storage.

Percentage of Space: HSM Level 1

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on HSM migration level 1 storage.

Percentage of Space: HSM Level 2

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on HSM migration level 2 storage.

Percentage of Space: HSM Backup

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on HSM backup storage. This total includes space used for multiple backup copies of the same data set.




Figure 4-42. Monthly Application Resource Usage / Plot

Monthly Application Resource Usage / Summary (SEGERE)

Source: VCADAA file at the MONTHS timespan
VCA_VS file at the MONTHS timespan
HSMmig file at the MONTHS timespan
HSBBAC file at the MONTHS timespan

Depending on your CA MICS installation options, these files may not be active in the MONTHS timespan. Check with your CA MICS Administrator if you have problems when trying to use this report.

Function: This facility allows the storage administrator to monitor key storage indicators on a month-by-month basis. These indicators enable accurate projections of future storage needs, both at the application level and the system level. The indicators reported by this facility are also useful for detecting trends that could indicate potential problems or areas in which storage resources are not being used efficiently. These reports, along with the Daily Application Resource Usage reports, provide both a long view and a short view of storage capacity and performance problems.

The ability to process data from multiple sources makes the reports produced even more valuable. Options allow the user to target a single storage media, such as real DASD or HSM backup storage, or to group data from multiple media on the same chart. This second type of reporting is useful when analyzing the flow of data across multiple Types of storage media.

The types of data that may be selected for reporting include the number of data sets, amount of space being allocated, percentages of data sets and space in certain categories, and the amount of allocated versus used space.

Features: Summarization options allow you to analyze the data based on its owner (using one or more of the CA MICS accounting fields), its SMS storage class, its SMS management class, or its location (such as real DASD, or HSM level 1 migration).

Data filtering allows you to include data from

selected summary groups. For example, you could choose to report on specific departments, groups or individuals, or on a specific SMS class.

Output reports from this facility include tabular reports, printer plots, or color graphic plots. When plots are requested, up to nine different elements may be plotted on the same chart.

SMS

Issues: The ability to summarize by an SMS storage or management class makes it easy to monitor the long-term activity within each of your defined classes. Many of the elements you can select for reporting are useful for monitoring an SMS migration effort, or to make sure SMS is operating correctly after migration is complete.

COLUMN DESCRIPTIONS

TITLE 4

The fourth report title line on each tabular report page indicates the indicator or element that is being reported on that page. Note the values displayed for each item will vary, depending on the input sources specified for the report. Any or all of the following values may appear:

Data Set Count: Total

Represents the total number of data sets found in the input data for this particular summary group.

Data Set Count: SMS-Managed

Represents the number of data sets found in the input data for this particular summary group that are managed by SMS.

Data Set Count: Unmanaged

Represents the number of data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Data Sets: Managed

Represents the percentage of all the data sets found in the input data for this particular summary group that are managed by SMS.

Percentage of Data Sets: Unmanaged

Represents the percentage of all the data sets found in the input data for this particular summary group that are not managed by SMS.

Space Allocated (MB): Total

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group.

Space Allocated (MB): Managed

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are managed by SMS.

Space Allocated (MB): Unmanaged

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Space: Managed

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are managed by SMS.

Percentage of Space: Unmanaged

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not managed by SMS.

Data Set Count: Non-VSAM

Represents the number of data sets found in the input data for this particular summary group that are not VSAM data sets.

Data Set Count: VSAM

Represents the number of data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Data Sets: Non-VSAM

Represents the percentage of data sets found in the input data for this particular summary group that are not VSAM data sets.

Percentage of Data Sets: VSAM

Represents the percentage of data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Space Used: All Data

Represents the percentage of space allocated versus used for all the data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the data sets in this group by the total tracks allocated. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Managed

Represents the percentage of space allocated versus used for the SMS-managed data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the managed data sets in this group by the total tracks allocated by the managed data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Unmanaged

Represents the percentage of space allocated versus used for the unmanaged data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the unmanaged data sets in this group by the total tracks allocated by the unmanaged data sets. Note that HSM data sets are not included in this calculation.

Space Allocated (MB): Non-VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not VSAM data sets.

Space Allocated (MB): VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Space: Non-VSAM

Based on the total space allocated by all the

data sets in this summary group, the percentage of that space belonging to data sets that are not VSAM data sets.

Percentage of Space: VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are VSAM data sets.

Percentage of Space Used: Non-VSAM

Represents the percentage of space allocated versus used for the non-VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the non-VSAM data sets in this group by the total tracks allocated by the non-VSAM data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: VSAM

Represents the percentage of space allocated versus used for the VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the VSAM data sets in this group by the total tracks allocated by the VSAM data sets. Note that HSM data sets are not included in this calculation.

Data Set Count: HSM Level 1

Represents the number of data sets found in the input data for this particular summary group that reside on HSM migration level 1 storage.

Data Set Count: HSM Level 2

Represents the number of data sets found in the input data for this particular summary group that reside on HSM migration level 2 storage.

Data Set Count: HSM Backup

Represents the number of data sets found in the input data for this particular summary group that reside on HSM backup storage. Note that multiple backup copies of the same data set do not affect this count, nor is any attempt made to indicate which of these data

sets may also reside on primary storage.

Space Allocated (MB): HSM Level 1

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on HSM migration level 1 storage.

Space Allocated (MB): HSM Level 2

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on HSM migration level 2 storage.

Space Allocated (MB): HSM Backup

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that reside on HSM backup storage. This total includes space used for multiple backup copies of the same data set.

Percentage of Data Sets: HSM Level 1

Represents the percentage of data sets found in the input data for this particular summary group that reside on HSM migration level 1 storage.

Percentage of Data Sets: HSM Level 2

Represents the percentage of data sets found in the input data for this particular summary group that reside on HSM migration level 2 storage.

Percentage of Data Sets: HSM Backup

Represents the percentage of data sets found in the input data for this particular summary group that reside on HSM backup storage. Note that multiple backup copies of the same data set do not affect this count, nor is any attempt made to indicate which of these data sets may also reside on primary storage.

Percentage of Space: HSM Level 1

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on HSM migration level 1 storage.

Percentage of Space: HSM Level 2

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on HSM migration level 2 storage.

Percentage of Space: HSM Backup

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that reside on HSM backup storage. This total includes space used for multiple backup copies of the same data set.

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Media

Storage Class

Management Class

Owner

The first column of the report indicates the summary group to which the totals on each line apply. Depending on the report option chosen, this summarization will be done by media (DASD, VSAM, MIGRATE, or BACKUP), SMS storage class, SMS management class, or based on the CA MICS accounting codes in the data. The label over the column will be one of the values shown above, depending on the summarization option chosen.

Value mmm yyyy

Indicates the value for this particular element for this particular summary group during the month and year indicated in the column heading.

Compound Growth Rate (%)

Reports the monthly compound growth rate (CGR) calculated for this group using the values from the months shown on this report. A value of 5.00, for example, indicates a growth rate of approximately 5% per month for this group, based on the monthly values included in this report.

MONTHLY APPLICATION RESOURCE USAGE (STGERE)						
CA MICS StorageMate						
ABC CORPORATION						
Data Set Count: Total						
----- SYSID=SYSA -----						
Owner	Value Nov 2006	Value Dec 2006	Value Jan 2007	Value Feb 2007	Value Mar 2007	Compound Growth Rate (%)
DEPT:100	836	966	1036	896	712	-3.93
DEPT:110	1973	2042	2200	2197	2231	3.12
DEPT:115	2328	2332	2655	2208	1522	-10.07
DEPT:120	19392	20955	21207	18845	11672	-11.92
DEPT:240	461	560	562	576	583	6.04
DEPT:250	170	191	229	514	545	33.81
DEPT:260A	52	50	53	57	56	2.23
DEPT:260B	.	39	474	401	711	162.41
DEPT:270	1521	1608	1645	1558	1223	-5.31
DEPT:280	39	52	71	75	69	15.37
DEPT:285	2	2	2	.	.	-8.23
DEPT:290	85	93	110	108	65	-6.45
DEPT:300	1367	1537	1693	1605	833	-11.64
DEPT:315	651	1038	733	678	601	-1.98
DEPT:400	1388	1441	1664	1431	590	-19.27
DEPT:430	833	1017	1274	1152	947	3.26
DEPT:430X	834	1069	1331	1521	1609	17.86
DEPT:500	79	77	95	110	101	6.29
DEPT:550	3350	3960	3667	3824	5394	12.65
DEPT:590	514	620	713	847	712	8.47
DEPT:600I	63	63	65	70	70	2.77
DEPT:799	23	24	22	22	23	0.14
DEPT:999	9	9	9	9	9	0.00
-----	-----	-----	-----	-----	-----	-----
SYSID	35968	39746	41511	38705	30279	
	=====	=====	=====	=====	=====	

Figure 4-43. Monthly Application Resource Usage / Summary

Monthly DASD Resource Usage / Plot (STGERH)

Source: VCAVOA file at the MONTHS timespan

Function: This facility allows the storage administrator to monitor key storage indicators on a month-by-month basis. These indicators enable the user to project long-term storage needs, both at the hardware unit level (volume, volume group, or device type) and the system level. The indicators reported by this facility are also useful for detecting problems with storage resources that are not being used efficiently. These reports, along with the Daily DASD Resource Usage reports, provide both a long view and a short view of storage capacity and performance problems, measured at the hardware level.

The Daily/Monthly DASD Resource Usage reports complement the Daily/Monthly Application Resource Usage reports to give a complete picture of storage usage and capacity. While the application reports tend to measure storage capacity and consumption at the application level, the DASD reports take a more hardware-oriented look at similar data. These reports are useful for tasks such as planning for new DASD, or balancing volumes between volume groups or SMS storage groups. Both sets of reports allow you to measure and control growth at the application and hardware levels.

The types of data that can be selected for reporting include the number of data sets, the amount of space being allocated, percentages of data sets and space in certain categories, the amount of allocated versus used space, and other measures of performance and capacity.

Features: Summarization options allow you to analyze the data based on volume serial number, device type, volume group or SMS storage group, or to provide totals for each system.

Data filtering allows you to include data from selected groups. For example, you can choose to report on specific volumes, device types, or SMS storage groups.

Output reports from this facility include tabular reports, printer plots, or color graphic plots. When plots are requested, up to nine different elements can be plotted on the same chart.

SMS

Issues: The ability to summarize by SMS storage groups makes it easy to monitor the long-term activity within each of your defined groups. Many of the elements you can select for reporting are useful for monitoring an SMS migration effort, or to make sure SMS is operating correctly after migration is complete.

ELEMENT DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Volume

Volume Group

Device Type

Unless summarization by SYSID only is chosen, the descriptive line at the top of each page will also indicate the summary group to which the values on this plot apply. Depending on the report option chosen, this summarization will be done by volume, volume group (SMS storage group), or device type. The label to the left of the summary value will be one of the values shown above, depending on the summarization option chosen.

Vertical Axis (Value)

The numeric values distributed vertically down the left side of each page represent the scale of the numeric values being plotted on that page. When requesting color graphics plots, you can specify the values that will appear on this axis. This should be done with care, because all plots and groups will have to adapt to the values you specify. For printer plots, or when values are not specified, SAS will generate the values automatically, based on the content of the plotted data. We suggest that these SAS defaults be used whenever possible, as they will avoid many potential problems. Care should also be taken when specifying multiple elements on one plot,

that the elements chosen will have a similar range of values.

Horizontal Axis (Year/Month)

Plots on the horizontal axis represent the month corresponding to the value being plotted. Each entry is presented in the format 'yy/mm'. The number of entries appearing on the horizontal axis is determined by the number of cycles of CA MICS data specified as input to the inquiry.

Footnotes

The footnotes section of each graph is important, as it indicates the names of the data element(s) that are plotted on the graph. Up to nine elements can be plotted on the same graph, in which case nine footnotes will appear at the bottom of each page. For printer plots, an alphabetic character ranging from A through I appear to the left of each description, identifying the character used on the plot to represent that particular element. On color graphic plots, the letter does not appear, as the color of the footnote should correspond with the color of the plot point or line on the graph. If different colors do not appear for each point and footnote, modify your MICF options to specify different colors. A brief description of the possible values in the footnote section, and what those values represent, is given below:

Data Set Count: Total

Represents the total number of data sets found in the input data for this particular summary group.

Data Set Count: SMS-Managed

Represents the number of data sets found in the input data for this particular summary group that are managed by SMS.

Data Set Count: Unmanaged

Represents the number of data sets found in the input data for this particular summary group that are not managed by SMS.

Data Set Count: Non-VSAM

Represents the number of data sets found in the input data for this particular summary

group that are not VSAM data sets.

Data Set Count: VSAM

Represents the number of data sets found in the input data for this particular summary group that are VSAM data sets.

Space Allocated (MB): Total

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group.

Space Allocated (MB): Managed

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are managed by SMS.

Space Allocated (MB): Unmanaged

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not managed by SMS.

Space Allocated (MB): Non-VSAM

Represents the total space allocated, in megabytes, by all the non-VSAM data sets found in the input data for this particular summary group.

Space Allocated (MB): VSAM

Represents the total space allocated, in megabytes, by all the VSAM data sets found in the input data for this particular summary group.

Percentage of Data Sets: Managed

Represents the percentage of all the data sets found in the input data for this particular summary group that are managed by SMS.

Percentage of Data Sets: Unmanaged

Represents the percentage of all the data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Data Sets: Non-VSAM

Represents the percentage of all the data sets found in the input data for this particular

summary group that are non-VSAM data sets.

Percentage of Data Sets: VSAM

Represents the percentage of all the data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Space: Managed

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are managed by SMS.

Percentage of Space: Unmanaged

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not managed by SMS.

Percentage of Space: Non-VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not VSAM.

Percentage of Space: VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are VSAM.

Percentage of Space Used: All Data

Represents the percentage of space allocated versus used for all the data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the data sets in this group by the total tracks allocated. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Non-VSAM

Represents the percentage of space allocated versus used for the non-VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the non-VSAM data sets in this group by the total tracks allocated by the non-VSAM data sets. Note that HSM data

sets are not included in this calculation.

Percentage of Space Used: VSAM

Represents the percentage of space allocated versus used for the VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the VSAM data sets in this group by the total tracks allocated by the VSAM data sets. Note that HSM data sets are not included in this calculation.

Total Storage Capacity (MB)

Represents the total published storage capacity for all the volumes in this particular summary group, expressed in megabytes.

Total Capacity Allocated (MB)

Represents the total amount of space within this particular summary group that has been allocated, and is no longer available for allocation by other data sets. This value is expressed in megabytes, and represents the space allocated by data sets and system areas such as VTOCs.

Total Capacity Free (MB)

Represents the total amount of space within this particular summary group that is available for allocation by other data sets. This value is expressed in megabytes.

Percentage of Capacity Allocated

Represents the percentage of space within this particular summary group that has been allocated, and is no longer available for allocation by other data sets. This value represents the space allocated by data sets and system areas such as VTOCs.

Percentage of Capacity Free

Represents the percentage of space within this particular summary group that is available for allocation by other data sets.

Average Fragmentation Index

Represents the average of all the fragmentation indexes for all the volumes in

this summary group. The fragmentation index is used as a measurement of space fragmentation on each volume, with a low number indicating little fragmentation, and higher numbers indicating greater amounts of fragmentation.

Free VTOC Index Records

Represents the total number of free VTOC index records for all of the volumes in this particular summary group. For volumes with indexed VTOCs, this number indicates the amount of index expansion space available.




Figure 4-44. Monthly DASD Resource Usage / Plot

Monthly DASD Resource Usage / Summary (STGERH)

Source: VCAVOA file at the MONTHS timespan

Function: This facility allows the storage administrator to monitor key storage indicators on a month-by-month basis. These indicators enable the user to project long-term storage needs, both at the hardware unit level (volume, volume group, or device type) and the system level. The indicators reported by this facility are also useful for detecting problems with storage resources that are not being used efficiently. These reports, along with the Daily DASD Resource Usage reports, provide both a long view and a short view of storage capacity and performance problems, measured at the hardware level.

The Daily/Monthly DASD Resource Usage reports complement the Daily/Monthly Application Resource Usage reports to give a complete picture of storage usage and capacity. While the application reports tend to measure storage capacity and consumption at the application level, the DASD reports take a more hardware-oriented look at similar data. These reports are useful for tasks such as planning for new DASD, or balancing volumes between volume groups or SMS storage groups. Both sets of reports allow you to measure and control growth at both the application and hardware levels.

The types of data that can be selected for reporting include the number of data sets, the amount of space being allocated, percentages of data sets and space in certain categories, the amount of allocated versus used space, and other measures of performance and capacity.

Features: Summarization options allow you to analyze the data based on volume serial number, device type, volume group or SMS storage group, or to provide totals for each system.

Data filtering allows you to include data from selected groups. For example, you can choose to report on specific volumes, device types, or SMS storage groups.

Output reports from this facility include tabular reports, printer plots, or color graphic plots. When plots are requested, up to nine different elements can be plotted on the same chart.

SMS

Issues: The ability to summarize by SMS storage groups makes it easy to monitor the long-term activity within each of your defined groups. Many of the elements you can select for reporting are useful for monitoring an SMS migration effort, or to make sure SMS is operating correctly after migration is complete.

COLUMN DESCRIPTIONS

TITLE 4

The fourth report title line on each tabular report page indicates the indicator or element that is being reported on that page. Note the values displayed for each item will vary, depending on the input sources specified for the report. Any or all of the following values may appear:

Data Set Count: Total

Represents the total number of data sets found in the input data for this particular summary group.

Data Set Count: SMS-Managed

Represents the number of data sets found in the input data for this particular summary group that are managed by SMS.

Data Set Count: Unmanaged

Represents the number of data sets found in the input data for this particular summary group that are not managed by SMS.

Data Set Count: Non-VSAM

Represents the number of data sets found in the input data for this particular summary group that are not VSAM data sets.

Data Set Count: VSAM

Represents the number of data sets found in the input data for this particular summary group that are VSAM data sets.

Space Allocated (MB): Total

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group.

Space Allocated (MB): Managed

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are managed by SMS.

Space Allocated (MB): Unmanaged

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not managed by SMS.

Space Allocated (MB): Non-VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are not VSAM data sets.

Space Allocated (MB): VSAM

Represents the total space allocated, in megabytes, by all the data sets found in the input data for this particular summary group that are VSAM data sets.

Percentage of Data Sets: Managed

Represents the percentage of all the data sets found in the input data for this particular summary group that are managed by SMS.

Percentage of Data Sets: Unmanaged

Represents the percentage of all the data sets found in the input data for this particular summary group that are not managed by SMS.

Percentage of Data Sets: Non-VSAM

Represents the percentage of all the data sets found in the input data for this particular summary group that are not VSAM.

Percentage of Data Sets: VSAM

Represents the percentage of all the data sets found in the input data for this particular summary group that are VSAM.

Percentage of Space: Managed

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are managed by SMS.

Percentage of Space: Unmanaged

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not managed by SMS.

Percentage of Space: Non-VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are not VSAM data sets.

Percentage of Space: VSAM

Based on the total space allocated by all the data sets in this summary group, the percentage of that space belonging to data sets that are VSAM data sets.

Percentage of Space Used: All Data

Represents the percentage of space allocated versus used for all the data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the data sets in this group by the total tracks allocated. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: Non-VSAM

Represents the percentage of space allocated versus used for the non-VSAM data sets found in the input data for this particular summary group. This is calculated by dividing the total tracks used by the non-VSAM data sets in this group by the total tracks allocated by the non-VSAM data sets. Note that HSM data sets are not included in this calculation.

Percentage of Space Used: VSAM

Represents the percentage of space allocated versus used for the VSAM data sets found in the input data for this particular summary

group. This is calculated by dividing the total tracks used by the VSAM data sets in this group by the total tracks allocated by the VSAM data sets. Note that DFHSM data sets are not included in this calculation, because they are always fully used.

Total Storage Capacity (MB)

Represents the total published storage capacity for all the volumes in this particular summary group, expressed in megabytes.

Total Capacity Allocated (MB)

Represents the total amount of space within this particular summary group that has been allocated, and is no longer available for allocation by other data sets. This value is expressed in megabytes, and represents the space allocated by data sets and system areas such as VTOCs.

Total Capacity Free (MB)

Represents the total amount of space within this particular summary group that is available for allocation by other data sets. This value is expressed in megabytes.

Percentage of Capacity Allocated

Represents the percentage of space within this particular summary group that has been allocated, and is no longer available for allocation by other data sets. This value represents the space allocated by data sets and system areas such as VTOCs.

Percentage of Capacity Free

Represents the percentage of space within this particular summary group that is available for allocation by other data sets.

Average Fragmentation Index

Represents the average of all the fragmentation indexes for all the volumes in this summary group. The fragmentation index is used as a measurement of space fragmentation on each volume, with a low number indicating little fragmentation, and higher numbers indicating greater amounts of

fragmentation.

Free VTOC Index Records

Represents the total number of free VTOC index records for all of the volumes in this particular summary group. For volumes with indexed VTOCs, this number indicates the amount of index expansion that may be done before running out of index space.

SYSID

Unless summarization by SYSID is chosen, this value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

SYSID

Volume

Volume Group

Device Type

The left-most data column on each report line indicates the summary group to which the values on that line apply. Depending on the report option chosen, this summarization will be done by volume, volume group (SMS storage group), device type, or system (SYSID). The label at the top of the column will be one of the values shown above, depending on the summarization option chosen.

Value mmm yyyy

Indicates the value for this particular element for this particular summary group during the month indicated in the column heading. The month is represented in the format 'mmm yyyy'.

Compound Growth Rate (%)

Reports the monthly compound growth rate (CGR) calculated for this group using the values from the months shown on this report. A value of 5.00, for example, indicates a growth rate of approximately 5% per month for this group, based on the monthly values included in this report.

MONTHLY DASD RESOURCE USAGE (STGERH)											
CA MICS StorageMate											
ABC CORPORATION											
Space Allocated (MB): Managed											
----- SYSID=IMS8 -----											
Volume Group	Value Mar 1991	Value Apr 1991	Value May 1991	Value Jun 1991	Value Jul 1991	Value Aug 1991	Value Sep 1991	Value Oct 1991	Value Nov 1991	Value Dec 1991	Compound Growth Rate (%)
LARGE	23446	19504	23594	26118	26561	29217	26771	29376	25415	25919	1.12
TECHSUP	7348	7001	7736	11415	10181	10221	10996	11344	11746	11881	5.48
TEMP	404	488	777	574	1372	1024	427	1205	1090	357	-1.36
TSO	31301	32379	36368	40962	43134	46017	44826	44992	40037	35683	1.47
SYSID	62499	59372	68475	79069	81247	86478	83020	86918	78289	73840	
	62499	59372	68475	79069	81247	86478	83020	86918	78289	73840	

Figure 4-45. Monthly DASD Resource Usage / Summary

New Data Set Allocation / Chart (STGECG)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSM MIG file at the DETAIL timespan
Tape library information (external)

Function: This facility allows the storage administrator to monitor the day-to-day allocation of new data sets. This permits early identification of groups or users that are allocating excessive space or using space incorrectly. One of the most common causes of space shortages in a typical installation is the intentional or unintentional allocation of large new data sets. These reports will help solve that type of problem, as new data sets occupying large amounts of space can be easily identified and corrected before they cause space shortages.

The facility permits you to indicate which data sets are to be considered new for the purposes of the reports. This option can help provide many different views of the same data. For example, you could analyze all the data sets created yesterday to isolate problem data sets, or you could look at all data sets created in the last month to establish growth patterns for each department.

Because this facility also supports input data describing the tapes in your tape library, the reports can also prove useful to those measuring the usage and growth of that resource. Those persons responsible for tuning DFHSM or DFSMS can also benefit from running these reports with HSM input to determine how soon new data sets are migrated to HSM.

Features: Summary reports can show new allocation trends based on volume serial number, volume group (or SMS storage group), system, SMS management class, SMS management class, one or more data set name nodes, one or more CA MICS accounting fields, or by using an external table of data set name patterns. Each summary report can be produced as a tabular report, printer graphic report, or color graphic report. Up to two different summary elements may be specified, to provide you with

many different data views.

If a summary report indicates a potential problem or anomaly, a detailed report can be produced to give more information about individual new data sets. You can specify whether all new data sets should be listed or just the largest. You can also specify whether the sort order of the detail report should be based on data set name or size.

Data filtering allows you to include data from selected groups. For example, you can choose to report on specific volumes, volume groups, or select new data sets allocated by only certain users or departments.

SMS

Issues: The reports in this facility will help you maximize your storage usage while SMS implementation is taking place. Even after SMS is fully implemented, poor allocation practices will still occur, as users can still override many of the SMS defaults. The ability to monitor new data set activity by storage group, storage class or management class will also help you make sure your SMS system is working as designed.

ELEMENT DESCRIPTIONS

SYSID

Management Class

Storage Class

Volume

Volume Group

Owner

If both a major and minor summarization element are selected, one of these labels will appear at the top of each chart. Based on the option you selected as the major summarization element, these values represent a System Identifier (SYSID), an SMS management class, an SMS storage class, a volume serial number, or a volume group (or SMS storage group). If the label "Owner" appears, this represents summarization based on one or more data set name nodes, one or more CA MICS accounting fields, or based on an external Data Set Group table.

Pie Slices

Each slice represents the amount of space, in megabytes, allocated by new data sets belonging to this particular summarization group.

Pie Labels

Each label identifies the summary group represented by this pie slice. The value that appears will depend on the major and minor summarization options selected. Depending on the option chosen, each label may represent either a SYSID, management class, storage class, volume serial number, volume group name (SMS storage group name), a portion of the data set name, a portion of the CA MICS accounting codes, or the result of a table search using a Data Set Group table. The example shown specified that the Volume Group should be the major summarization element, and the minor summarization element should be the output of a Data Set Group table. The result of this execution would produce one chart for each Volume Group, with each pie slice on the chart representing new space allocated in that volume group by various users. The names of the users are determined by matching the data set names of new data sets with the specified Data Set Group table.




Figure 4-46. New Data Set Allocation / Chart

New Data Set Allocation / Summary (STGECG)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
Tape library information (external)

Function: This facility allows the storage administrator to monitor the day-to-day allocation of new data sets. This permits early identification of groups or users that are allocating excessive space or using space incorrectly. One of the most common causes of space shortages in a typical installation is the intentional or unintentional allocation of large new data sets. These reports will help solve that type of problem, as new data sets occupying large amounts of space can be easily identified and corrected before they cause space shortages.

The facility permits you to indicate which data sets are to be considered new for the purposes of the reports. This option can help provide many different views of the same data. For example, you could analyze all the data sets created yesterday to isolate problem data sets, or you could look at all data sets created in the last month to establish growth patterns for each department.

Because this facility also supports input data describing the tapes in your tape library, the reports can also prove useful to those measuring the usage and growth of that resource. Those persons responsible for tuning DFHSM or DFSMS can also benefit from running these reports with HSM input to determine how soon new data sets are migrated to HSM.

Features: Summary reports can show new allocation trends based on volume serial number, volume group (or SMS storage group), system, SMS management class, SMS management class, one or more data set name nodes, one or more CA MICS accounting fields, or by using an external table of data set name patterns. Each summary report can be produced as a tabular report, printer graphic report, or color graphic report. Up to two different summary elements may be specified, to provide you with

many different data views.

If a summary report indicates a potential problem or anomaly, a detailed report can be produced to give more information about individual new data sets. You can specify whether all new data sets should be listed or just the largest. You can also specify whether the sort order of the detail report should be based on data set name or size.

Data filtering allows you to include data from selected groups. For example, you can choose to report on specific volumes, volume groups, or select new data sets allocated by only certain users or departments.

SMS

Issues: The reports in this facility will help you maximize your storage usage while SMS implementation is taking place. Even after SMS is fully implemented, poor allocation practices will still occur, as users can still override many of the SMS defaults. The ability to monitor new data set activity by storage group, storage class or management class will also help you make sure your SMS system is working as designed.

COLUMN DESCRIPTIONS

SYSID

Management Class

Storage Class

Volume

Volume Group

Owner

This report requires at least one summarization level, with a second one being optional. If two levels are selected, the major level will appear at the top of each page, and the minor level will appear as the left-most column in each report. If only one summarization level was specified, it will appear as the left-most column. The description for each summary element will be one of those values shown above. These values represent summarization by System Identifier (SYSID), SMS management class, SMS storage class, volume serial number, or volume group (or SMS storage group). If the label "Owner" appears, this represents summarization based on one or more

data set name nodes, one or more CA MICS accounting fields, or based on an external Data Set Group table.

New Data Sets

The number of new data sets found in the input data belonging to this summary group. The age of a data set that is considered "new" is determined by the options you coded on the report options panel.

% Data Sets New

The percentage of all data sets found in the input data for this summary group that represent new data sets. If 100 data sets were found for this group and 25 of them were new, this value would be 25%.

New Space (MB)

The amount of storage space, expressed in megabytes, allocated by the new data sets found for this summary group.

% Space New

The percentage of all allocated space found in the input data for this summary group that represent space allocated by new data sets. If all data sets in this group allocated 100 megabytes of space, and 40 of those megabytes were allocated by new data sets, this value would be 40%.

Avg New Size (MB)

The average amount of space allocated, expressed in megabytes, by the new data sets in this summary group. This is calculated by dividing the total amount of space allocated by new data sets by the number of new data sets.

Largest New Data Set Information:

All of the items on the right side of this report refer to the largest data set found among all the new data sets in this summary group.

Name

Specifies the data set name of the largest new data set in this group.

Space Alloc (MB)

Specifies the amount of space allocated, in

megabytes, by the largest new data set in this summary group.

Volume Serial Number

Specifies the volume on which the largest new data set in this summary group resides.

NEW DATA SET ALLOCATION (STGECG)									
CA MICS StorageMate									
ABC CORPORATION									
Input Contained Data Sets Created: 30DEC91 -- 01JAN92									
----- Volume Group=LOGGRP -----									
Owner	New Data Sets	% Data Sets New	New Space (MB)	% Space New	Avg New Size (MB)	Largest New Data Set Information: Name	Space Alloc (MB)	Volume Serial Number	
SALES	70	17.5	2629.9	19.8	37.6	CAPTEST.I398834B.DUMP	618.8	LOG001	
MARKETING	6	9.7	86.7	7.6	14.4	SYSLOG.IMS8.DAT91364	17.8	LOG003	
PAYROLL	3	4.9	18.1	4.7	6.0	JOBLOG.IMS8.DAT92001	8.5	LOG003	
DATA-CENTER	1	1.2	12.3	0.4	12.3	PRODTTEST.CLONE.NETWRK41	12.3	LOG015	
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
OWNER	80		2747.0						
----- Volume Group=TSTGRP -----									
Owner	New Data Sets	% Data Sets New	New Space (MB)	% Space New	Avg New Size (MB)	Largest New Data Set Information: Name	Space Alloc (MB)	Volume Serial Number	
SALES	162	2.8	360.4	2.8	2.2	DEVKB00.AA01.DISTLIST	54.4	TST031	
DATA-CENTER	39	2.7	309.5	5.7	7.9	MSR.MR.ISPTLIB	63.7	TST057	
MARKETING	26	3.5	61.5	2.2	2.4	TSOLOG.SMPLOG	8.5	TST057	
FINANCIAL	18	4.0	2.5	0.1	0.1	QTST.INFOREP.MPCSINOP	0.5	TST033	
CORPORATE	2	1.9	1.7	0.5	0.8	FLEX.REV1.LINKLOG	0.8	TST028	
PAYROLL	6	1.7	0.5	0.0	0.1	JOBLOG.IMS8.DAT91365	0.1	TST056	
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
OWNER	253		736.1						

Figure 4-47. New Data Set Allocation / Summary

New Data Set Allocation / Detail (STGECG)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
Tape library information (external)

Function: This facility allows the storage administrator to monitor the day-to-day allocation of new data sets. This permits early identification of groups or users that are allocating excessive space or using space incorrectly. One of the most common causes of space shortages in a typical installation is the intentional or unintentional allocation of large new data sets. These reports will help solve that type of problem, as new data sets occupying large amounts of space can be easily identified and corrected before they cause space shortages.

The facility permits you to indicate which data sets are to be considered new for the purposes of the reports. This option can help provide many different views of the same data. For example, you could analyze all the data sets created yesterday to isolate problem data sets, or you could look at all data sets created in the last month to establish growth patterns for each department.

Because this facility also supports input data describing the tapes in your tape library, the reports can also prove useful to those measuring the usage and growth of that resource. Those persons responsible for tuning DFHSM or DFSMS can also benefit from running these reports with HSM input to determine how soon new data sets are migrated to HSM.

Features: Summary reports can show new allocation trends based on volume serial number, volume group (or SMS storage group), system, SMS management class, SMS management class, one or more data set name nodes, one or more CA MICS accounting fields, or by using an external table of data set name patterns. Each summary report can be produced as a tabular report, printer graphic report, or color graphic report. Up to two different summary elements may be specified, to provide you with many different data views.

If a summary report indicates a potential problem or anomaly, a detailed report can be produced to give more information about individual new data sets. You can specify whether all new data sets should be listed or just the largest. You can also specify whether the sort order of the detail report should be based on data set name or size.

Data filtering allows you to include data from selected groups. For example, you can choose to report on specific volumes, volume groups, or select new data sets allocated by only certain users or departments.

SMS

Issues: The reports in this facility will help you maximize your storage usage while SMS implementation is taking place. Even after SMS is fully implemented, poor allocation practices will still occur, as users can still override many of the SMS defaults. The ability to monitor new data set activity by storage group, storage class or management class will also help you make sure your SMS system is working as designed.

COLUMN DESCRIPTIONS

SYSID

This value appears at the top of each page and identifies the System Identifier (system) to which this data applies.

Data Set Name

Indicates the name of the new data set. All of the data sets that appear on this detailed report are considered "new" as determined by the options you coded on the report options panel.

Alloc Space (MB)

The amount of storage space being allocated by this data set, expressed in megabytes.

Used Space (MB)

The amount of storage space allocated by this data set that is actually being used for the storage of data. The value is expressed in megabytes.

Volume Serial Number

The serial number of the volume on which this data set resides.

Date Created

The date this data set was created, in the format ddmmmyy.

Date Last Used

The date this data set was last referenced, in the format ddmmmyy.

Data Set Org.

The data set organization for this data set.

Block Size

The block size (BLKSIZE) of this data set.

Record Length

The logical record length (LRECL) of this data set.

Record Format

The record format (RECFM) of this data set.

NEW DATA SET ALLOCATION (STGECG)									
CA MICS StorageMate - DETAIL									
ABC CORPORATION									
Input Contained Data Sets Created: 30DEC91 -- 01JAN92									
----- SYSID=IMS8 -----									
Data Set Name	Alloc Space (MB)	Used Space (MB)	Volume Serial Number	Date Created	Date Last Used	Data Set Org.	Block Size	Record Length	Record Format
CAPTEST.I398834B.DUMP	618.75	618.19	LOG001	31DEC91	31DEC91	PS	4160	4160	F
DEV001.MULTIVOL.TESTDATA	273.31	273.31	TST032	30DEC91	30DEC91	PS	27998	4096	VB
CAPTEST.I400913.DUMP	117.29	117.23	LOG009	31DEC91	31DEC91	PS	4160	4160	F
DEVTP11.#TPP.D398244.R20	99.45	99.33	LOG003	31DEC91	31DEC91	PS	1310	1310	F
DEVTP11.#TPP.D398244.R202	81.59	81.37	LOG015	31DEC91	31DEC91	PS	1310	131	FB
DEVTP11.#TPP.D400445.R201	75.64	75.02	LOG005	30DEC91	30DEC91	PS	4160	4160	F
DEVTP11.#TPP.D403055.R201	68.00	67.20	LOG005	31DEC91	31DEC91	PS	4160	4160	FB
DEVTP11.#TPP.D403055.R202	67.14	66.35	LOG015	31DEC91	31DEC91	PS	4160	4160	FB
DEVTP11.#TPP.D398244.R203	65.45	64.65	LOG014	31DEC91	31DEC91	PS	1310	131	FB
DEVTP11.#TPP.D401608.R201	64.59	64.43	LOG011	30DEC91	30DEC91	PS	4160	4160	FB
MSR.MX.ISPTLIB	63.75	0.62	TST057	31DEC91	.	P0	6160	80	FB
MSR.MX.SOURCE	63.75	34.79	TST033	31DEC91	31DEC91	P0	6160	80	FB
DEVKB11.AA01.DISTLIST	54.39	53.26	TST031	30DEC91	01JAN92	P0	6050	121	FBM
DEV118.LAB.ALL.DETAIL	51.00	0.11	LOG005	31DEC91	31DEC91	DA	32760	32756	U
DEV118.LAB.ONLY.DETAIL	51.00	0.11	LOG003	31DEC91	31DEC91	DA	32760	32756	U

SYSID	1815.10	1615.98							
	=====	=====							
	1815.10	1615.98							
N=	15								

Figure 4-48. New Data Set Allocation / Detail

Non-VSAM Attributes / Summary (STGEFB)

Source: VCADAA file at the DETAIL timespan
Tape library information (external)

Function: Identifies common patterns of data set attributes within your non-VSAM data sets. This is useful in identifying data sets with similar attributes that may have similar processing requirements.

This report also helps identify groups of data sets that have been allocated with attributes that may cause performance problems or poor space utilization.

Features: Users can choose to ignore data set block size when grouping, as this may not be important when using automated software that does automatic reblocking.

Users can also choose to drop those groups from the report that represent a small percentage of their total data sets.

Generation of the summary report is automatic.

SMS

Issues: After running this summary report, three additional reports that expand on the information presented in this summary can be requested.

SMS allows similar data sets to be managed under one Data Class, and this report helps you identify and define those classes.

You can easily determine the most common groups of data sets, thus allowing data sets in those groups to be automated quickly. The planning information provided by the report allows group definitions to be made with less chance of error.

COLUMN DESCRIPTIONS

Report Group

Represents a group of data sets that have identical attributes, as defined by the columns 'DSORG' through 'Retention Period' on each report line. Report Group is important, as it is used in all other Non-VSAM Attributes reports to refer to

this common group of data sets. The number of report groups is determined by the characteristics of the data analyzed, as well as by values entered on the report option screen.

Number of Data Sets

The number of data sets within the data analyzed that contained attributes identical to those defined by this group.

Percent of All Data Sets

The percentage of data sets within the data analyzed that were assigned to this report group. Note that any data sets excluded by filtering values on the report option panel are not included when calculating this percentage.

DSORG

The organization of the data within the data set. Note that all the data sets assigned to this report group have the data set organization listed under this column. A value of 'PS' will always appear for tape data sets.

BLKSIZE

The maximum length of a physical block written into this data set. Note that all the data sets assigned to this report group have the block size listed under this column. A value on the report option screen allows this value to be set to zero for all data sets, so that block size will not be a factor when grouping. If all values under this column are zero, that option was probably used.

LRECL

The maximum length of each logical record in this data set. Note that all the data sets assigned to this report group have the logical record length listed under this column.

RECFM

Note that all the data sets assigned to this report group have the record format listed under this column. If the option was chosen on the report options screen to ignore block size when grouping, all 'B' (Blocked) values will be removed from RECFM, and will not be listed on this report.

Relative Key Position

The relative position of the first byte of the key within the logical record for ISAM files. Note that all the data sets assigned to this report group have the same relative key position listed under this column. This value will always be zero for tape data sets.

Key Length

The length of the key associated with each data block in a keyed data set. Note that all the data sets assigned to this report group have the same key length listed under this column. This value will always be zero for tape data sets.

Retention Period

The retention period specified for all the data sets assigned to this report group. While retention period is not reported directly in the input data, it can be computed by subtracting the date the data set was created from its expiration date. If no expiration date was specified, the retention period will be set to zero. Note that you have the option to set this to zero when using tape data sets.

Average Data Set Size (KB)

The average size, expressed in thousands of bytes (kilobytes), of all data sets in this reporting group. This is calculated by taking the total amount of tracks allocated by all data sets in this group, and dividing it by the number of data sets. This value can be useful when trying to set a default amount of space that could be assigned to the data sets in this group.

Minimum Data Set Size (KB)

The smallest amount of space allocated, expressed in thousands of bytes (kilobytes), detected for all of the data sets in this reporting group. A value of zero indicates that at least one data set in the group is allocated with zero extents and is, therefore, using no space for that data set. This value can be useful when trying to set a default amount of space that could be assigned to the data sets in this group.

Maximum Data Set Size (KB)

The largest amount of space allocated, expressed in thousands of bytes (kilobytes), detected for

all of the data sets in this reporting group.
This value can be useful when trying to set a
default amount of space that could be assigned to
the data sets in this group.

NON-VSAM ATTRIBUTES (STGEFB)												1
CA MICS StorageMate - SUMMARY												
ABC CORPORATION												
Report Group	Number of Data Sets	Percent of All Data Sets	DSORG	BLKSIZE	LRECL	RECFM	Relative Key Position	Key Length	Retention Period	Average Data Set Size(KB)	Minimum Data Set Size(KB)	Maximum Data Set Size(KB)
1	2936	38.9	PS	0	2048	FS	0	0	0	803	57	67543
2	2333	30.9	P0	0	80	F	0	0	0	5822	47	1068188
3	1162	15.4	P0	0	0	U	0	0	0	8304	47	1393063
4	268	3.5	PS	0	80	F	0	0	0	6526	0	730938

Figure 4-49. Non-VSAM Attributes / Summary

Non-VSAM Attributes / Chart (STGEFB)

Source: VCADAA file at the DETAIL timespan
Tape library information (external)

Function: Expands on the information in the summary report, identifying the range of space being allocated by the non-VSAM data sets within an identified group.

Once you have identified groups of data sets with common attributes, it is often useful to determine the amount of space they represent, and the range of that space allocation.

Features: Identification of size patterns within common data set groups can help when calculating current utilization or planning for future growth.

This report can also help identify exceptional data sets that are improperly sized for the function they perform.

Color graphics is supported.

Generation of this report is optional, as only the summary report is produced automatically.

SMS

Issues: Just as the summary report helps identify potential for SMS data classes, this report helps you determine the default amount of space that should be allocated by each of those classes. This will allow your data classes to be established faster, and with less chance of error.

COLUMN DESCRIPTIONS

Report Group

Identifies the reporting group for which this chart is produced. Refer back to the summary report to determine the group of data sets that comprise this report group. A separate chart is produced for each summary group, so that allocation trends for common groups of data sets can be identified.

Vertical Axis

The vertical axis of the chart represents a list of allocation values ranging from 100 kilobytes

(thousand bytes) to 50000 kilobytes.

Horizontal Axis

The horizontal axis represents the cumulative number of data sets found within each allocation range. As this is a cumulative value, the chart line should grow longer as it moves down the page, representing almost all the data sets analyzed as it reaches the bottom.

Other

The numeric values to the right of each chart line represent the number of data sets in each range, the cumulative number, the percent of all data sets within each range, and the cumulative percentage.




Figure 4-50. Non-VSAM Attributes / Chart

Non-VSAM Attributes / Node Analysis (STGEFB)

Source: VCADAA file at the DETAIL timespan
Tape library information (external)

Function: Expands on the information in the summary report, providing a data set name node summary of the non-VSAM data sets within each group.

This is useful when the user needs to know the name structures of data sets with certain attributes.

Features: Report options allow analysis of up to nine nodes, (1st, 2nd, ...9th), a range of nodes, or the last node of the data set name.

You can specify that only the most commonly occurring nodes should appear on the report.

Once the summary report has identified a group of data sets with similar attributes, this report helps identify users or groups that own the data.

SMS

Issues: Establishing SMS data classes usually requires two steps: defining the attributes of the class and finding a way to assign each new data set to the correct class. This report helps with the second part of this process, as it identifies patterns that can be used in ACS routines to make class assignments.

COLUMN DESCRIPTIONS

Report Group

Identifies the reporting group for which this analysis report is produced. Refer back to the summary report to determine the group of data sets that comprise this report group. All data sets assigned to this particular report group will have their data set node(s) analyzed by this report.

Data Set Name Pattern (Node)

The data set node(s) represented by this summarization. The contents of this column are determined by the values coded on the report option screen for this report, and by the values found in your installation's data set names for

this report group. The occurrence of leading, trailing, or duplicate periods in this value indicates that the node(s) you specified were not present for the data set names represented by this group. For example, trying to analyze nodes 2-3 of data set name SYS1.PROCLIB would result in a value of 'PROCLIB.', while trying to use nodes 3-5 would result in a value of '..'.

Data Set Count

The number of data set names included in this report group that matched the data set name structure identified in the Data Set Name Pattern column. The sum of this column may not equal the number of data sets in the report group, as very small groups of common node(s) can be dropped from the analysis, as specified on the report options screen.

% of Group Total

The percentage of all the data set names included in this report group that matched the data set name structure identified in the Data Set Name Pattern column. The sum of this column may not equal 100%, as very small groups of common node(s) can be dropped from the analysis, as specified on the report options screen.

Average Size (KB)

The average size of each data set in this report group with this data set name pattern, expressed in thousands of bytes (kilobytes). This is calculated by adding the size of each data set in the group, and dividing by the number of data sets.

Minimum Size (KB)

The size of the smallest data set in this report group with this data set name pattern, expressed in thousands of bytes (kilobytes).

Maximum Size (KB)

The size of the largest data set in this report group with this data set name pattern, expressed in thousands of bytes (kilobytes).

Percentage of Data Sets Less Than:

This represents a cumulative percentage distribution line for all data sets within this

data set name pattern and this report group.
 Values along the line range from 100 kilobytes to
 10 megabytes, and represent the cumulative
 percentage of data sets having sizes within each
 range.

% Other

The percentage of data sets within this data set
 name pattern and this report group that are larger
 than 10 megabytes in size.

NON-VSAM ATTRIBUTES (STGEFB)																	30
CA MICS StorageMate - NODE ANALYSIS																	
ABC CORPORATION																	
----- Report Group (See Summary Report)= 1 -----																	
Data Set Name Pattern (Node)	Data Set Count	% of Group Total	Average Size (KB)	Minimum Size (KB)	Maximum Size (KB)	+-----	Percentage of Data	Sets	Less Than:	-----+	% Other						
						100 KB	250 KB	500 KB	750 KB	1 MB	2 MB	4 MB	6 MB	8 MB	10 MB		
I0211	224	7.63	1022	57	67543	50	69	77	84	86	95	97	97	97	97	3	
H9271	160	5.45	332	57	3570	92	92	92	92	92	92	100	100	100	100	0	
H9270	106	3.61	285	57	3570	93	93	93	93	93	93	100	100	100	100	0	
I0171	83	2.83	345	57	8556	75	83	87	92	92	96	99	99	99	100	0	
H9305	78	2.66	343	57	6516	82	87	88	91	92	94	99	99	100	100	0	
I0212	73	2.49	882	57	21305	51	70	77	86	88	96	96	96	96	96	4	
H9024	61	2.08	58	57	113	98	100	100	100	100	100	100	100	100	100	0	
I0061	58	1.98	263	57	4476	66	76	79	98	98	98	98	100	100	100	0	
I0040	56	1.91	522	57	8500	55	77	86	89	93	96	96	96	96	100	0	
I0058	56	1.91	985	57	24422	48	70	82	91	93	93	96	96	96	96	4	
I0030	47	1.60	342	57	4023	70	81	89	91	91	96	98	100	100	100	0	
I0031	46	1.57	623	57	18641	76	83	91	91	91	96	98	98	98	98	2	
I0009	44	1.50	287	57	7253	82	89	93	93	95	98	98	98	100	100	0	
I0205	42	1.43	1560	57	17055	48	62	71	83	83	86	88	93	93	93	7	
I0208	37	1.26	639	57	10369	51	76	81	89	89	92	95	97	97	97	3	
I0029	36	1.23	320	57	3286	64	83	89	89	89	97	100	100	100	100	0	
I0073	36	1.23	252	57	4193	78	86	94	94	94	97	97	100	100	100	0	
H9157	35	1.19	1737	57	33031	63	77	77	83	83	86	89	91	97	97	3	
I0197	34	1.16	1152	57	7026	41	59	65	71	74	79	94	97	100	100	0	
H9290	33	1.12	206	57	1983	58	82	91	97	97	100	100	100	100	100	0	
I0201	32	1.09	1054	57	11786	50	66	78	81	81	91	94	94	94	94	6	
TSMLINK	15	1.29	2750	1424	5100	0	0	0	0	0	60	67	100	100	100	0	
CMDLIB	14	1.20	4884	567	14166	0	0	0	14	14	21	21	93	93	93	7	
F4	13	1.12	20425	567	113324	0	0	0	8	23	46	77	77	77	77	23	
IMAGELIB	13	1.12	12152	6694	16996	0	0	0	0	0	0	0	0	38	54	46	
MLPALIB	12	1.03	7437	4250	12749	0	0	0	0	0	0	0	50	50	75	25	
PSBLIB	12	1.03	2300	57	8500	8	8	25	25	33	58	83	92	92	100	0	
SGIMLMD0	12	1.03	4659	1813	7649	0	0	0	0	0	17	17	83	100	100	0	

Figure 4-51. Non-VSAM Attributes / Node Analysis

Non-VSAM Attributes / Detail (STGEFB)

Source: VCADAA file at the DETAIL timespan
Tape library information (external)

Function: Expands on the information in the summary report, providing a detailed listing of the non-VSAM data sets within selected groups.

This is useful when you need to have detailed information about all the data sets allocated with a certain set of attributes.

Features: Used in conjunction with the summary report, this report identifies individual data sets that are allocated with attributes causing poor performance or poor space utilization.

Because this report can produce a large volume of output, users can specify that detailed information be provided for selected groups, for all groups, or for none of the groups.

SMS

Issues: Establishing SMS data classes usually requires two steps: defining the attributes of the class and finding a way to assign each new data set to the correct class. This report helps with the second part of this process, as it identifies patterns that can be used in ACS routines to make class assignments.

COLUMN DESCRIPTIONS

Report Group

Identifies the reporting group for which this detail report is produced. Refer back to the summary report to determine the group of data sets that comprise this report group. All data sets assigned to this particular report group will be listed on this detail report. Detailed reports of selected report groups can be requested from the report options screen.

Data Set Name

The name of the data set.

SYSID

The translated system identification that was

derived from ORGSYSID by the CA MICS SYSID Translation process.

DSORG

The organization of the data within the data set and if the data set contains any location-dependent information that would make it unmoveable. A value of 'PS' is always shown for tape data sets.

BLKSIZE

The maximum length of a physical block written into this data set. If this contains a value of zeros for all data sets, the option was entered on the report options screen that causes block size to be ignored when doing grouping.

LRECL

The maximum length of each logical record in this data set.

RECFM

If the option was chosen on the report options screen to ignore block size when grouping, all 'B' (Blocked) values will be removed from RECFM, and will not be listed on this report.

Rel. Key Pos.

The relative position of the first byte of the key within the logical record.

Key Len.

The length of the key associated with each data block in a keyed data set. This value is always zero for tape data sets.

Date Created

The date that the data set was first allocated to a DASD volume.

Ret. Days

The number of days specified for the data set as its retention period. This is calculated by subtracting the expiration date from the date the data set was created. A value of zero indicates no retention value was specified. Note that you have the option of always setting this to zero for tape data sets.

Date Last Used

The date that the data set was last referenced.

Space Alloc (KB)

The total amount of space allocated by this data set, expressed in thousands of bytes (kilobytes). This is calculated by multiplying the number of tracks allocated by the track capacity in kilobytes.

Space Used (KB)

The total amount of space used by this data set, expressed in thousands of bytes (kilobytes). This is calculated by multiplying the number of tracks used by the track capacity in kilobytes. This will always be the same value as space allocated for tape data sets.

NON-VSAM ATTRIBUTES (STGEFB)												6
CA MICS StorageMate - DETAIL												
ABC CORPORATION												
----- Report Group (Refer to Summary Report)=3 -----												
Data Set Name	SYSID	DSORG	BLKSIZE	LRECL	RECFM	Rel. Key Pos.	Key Len.	Date Created	Ret. Days	Date Last Used	Space Alloc (KB)	Space Used (KB)
CMC.CICS.CIC00T.LOADLIB3	SYSA	PO	0	0	U	0	0	17NOV89	0	27JUL90	4178	3656
CMC.CICS.LOADLIB	SYSA	PO	0	0	U	0	0	14NOV89	0	26JUL90	6290	6290
CMC.CICS.LOADLIB1	SYSA	PO	0	0	U	0	0	30JAN89	0	26JUL90	340	227
CMC.CICS.LOADLIB2	SYSA	PO	0	0	U	0	0	19DEC89	0	26JUL90	8403	7359
CMC.CICS.LOADLIB3	SYSA	PO	0	0	U	0	0	14NOV89	0	30JUL90	283	283
CMC.CICS.V1R7M1.DLOADLIB	SYSA	PO	0	0	U	0	0	18JAN89	0	22MAY90	4273	3893
CMC.CICS.V1R7M1.IP01.LINKLIB	SYSA	PO	0	0	U	0	0	13JAN89	0	21MAY90	712	95
CMC.CICS.V1R7M1.LOADLIB	SYSA	PO	0	0	U	0	0	13JAN89	0	30JUL90	5697	5507
CMC.CICS.V1R7M1.LOADLIB1	SYSA	PO	0	0	U	0	0	13JAN89	0	27JUN90	712	190
CMC.CICS.V1R7M1.LOADLIB2	SYSA	PO	0	0	U	0	0	13JAN89	0	21MAY90	712	47
CMC.CICS.V1R7M1.LOADLIB2.DBH	SYSA	PO	0	0	U	0	0	19SEP89	0	21MAY90	712	95
CMC.CICS.V1R7M1.TMPLD	SYSA	PO	0	0	U	0	0	08FEB89	0	21MAY90	1424	237
CMC.CICS.V2R1M0.DLOADLIB	SYSA	PO	0	0	U	0	0	12OCT89	0	22MAY90	4273	4273
CMC.CICS.V2R1M0.LOADLIB	SYSA	PO	0	0	U	0	0	12OCT89	0	26JUN90	6409	5839
CMC.CICS.V2R1M0.LOADLIB1	SYSA	PO	0	0	U	0	0	12OCT89	0	22MAY90	712	332
CMC.CICS.V2R1M0.LOADLIB2	SYSA	PO	0	0	U	0	0	12OCT89	0	22MAY90	712	47
CMC.ONLINE.V450.RPL	SYSA	PO	0	0	U	0	0	27APR90	0	26JUL90	427	190
CIC00P.LOADLIB1.NCC00211.MS000PE1	SYSA	PO	0	0	U	0	0	28JUN90	0	24JUL90	1424	237
CIC00T.DADSPLUS.V2R1.LOADLIB	SYSA	PO	0	0	U	0	0	14NOV89	0	28NOV89	2848	2563
CIC00T.ONLINES0.INTRTEST.LOADLIB	SYSA	PO	0	0	U	0	0	22MAR90	0	26JUL90	1852	1377
CIC00T.ONLINES0.REPORTS.LOADLIB	SYSA	PO	0	0	U	0	0	28FEB90	0	26JUL90	3513	2801
CIC00T.ONLINES0.PLDS	SYSA	PO	0	0	U	0	0	28FEB90	0	27APR90	380	190
CIC00T.ONLINES0.SUPEROPT.LOADLIB	SYSA	PO	0	0	U	0	0	28FEB90	0	27APR90	2753	2516
CIC00T.ONLINES0.VERIFY.LOADLIB	SYSA	PO	0	0	U	0	0	22MAR90	0	26JUL90	760	427
CIS00T.CLINKLIB.NCI00211.MS000PE1	SYSA	PO	0	0	U	0	0	17NOV89	0	13JUL90	22098	21813
DBA00T.PSBLIB00.NLN00330.MS000PE1	SYSA	PO	0	0	U	0	0	17NOV89	0	30JUL90	5723	1643
DCI.BOSTON.GHR4150.LOAD	SYSA	PO	0	0	U	0	0	31JUL90	0	31JUL90	850	170
DJH8940.XLIB.LOAD	SYSA	PO	0	0	U	0	0	31MAR89	0	26JUL90	332	332
DST9738.TICOLD.CMDLIB	SYSA	PO	0	0	U	0	0	31JUL90	0	31JUL90	712	332
DPS00T.CLINKLIB.NDP00621.MS000PE1	SYSA	PO	0	0	U	0	0	17NOV89	0	13JUL90	20398	16430
DPS00T.LINKLIB0.NDP00621.MS000PE1	SYSA	PO	0	0	U	0	0	17NOV89	0	12JUL90	20176	19605
DPS00T.LINKLIB0.NDP00621.MS000PE1	SYSA	PO	0	0	U	0	0	17NOV89	0	17NOV89	19660	170
DRC8732.SYS2.LINK	SYSA	PO	0	0	U	0	0	27JUL89	0	31JUL90	283	227
FC0.DBN.R4M0.DBNLOAD	SYSA	PO	0	0	U	0	0	04MAY90	0	31MAY90	2437	1870
DBT.V1R1M1.AFABMOD0	SYSA	PO	0	0	U	0	0	28AUG89	0	20MAR90	850	453
DBT.V1R1M1.HJZ1102.F2	SYSA	PO	0	0	U	0	0	28AUG89	0	20MAR90	227	227
DBT.V1R1M1.JJZ1111.F2	SYSA	PO	0	0	U	0	0	28AUG89	0	20MAR90	397	397
-----											9649585	5013422
GROUP											=====	=====
											9649585	5013422

Figure 4-52. Non-VSAM Attributes / Detail

Non-VSAM Wasted Space / Chart (STGECC)

Source: VCADAA file at the DETAIL timespan

Function: This chart helps to identify departments or groups that are wasting DASD storage space due to poor allocation practices. These problems can be caused either by allocating more space than is necessary, or by using an inefficient block size. The Storage Administrator can analyze either or both of these causes with the same chart. Correction of the problems detected by this report may recover a large amount of storage space with a small amount of effort.

Features: This chart could be run regularly to monitor wasted space, with more detailed reporting occurring when problems are detected. A wide variety of filtering and reporting options allow many views of the same data to be obtained. Both color graphics and printer graphics are supported for this chart.

SMS

Issues: Users planning for SMS could identify those groups with the worst space usage values, as they may be the best candidates for the immediate benefits of SMS. As groups migrate to SMS control, regular use of this report would verify that the expected space savings from SMS are being realized.

COLUMN DESCRIPTIONS

Pie Slices

The pie slices represent the total amount of space being wasted by the data sets in each summarization group, in millions of bytes (megabytes). Depending on the values coded on the report option screen, this can represent space wasted due to poor block sizes, space wasted due to over-allocation, or both.

Pie Labels

The labels on each slice identify the summarization key for this chart. Depending on the values specified on the report option screen, this will represent one or more VCA account elements, one or more nodes of the data set name, or will contain a value derived from your Data Set

Group table.




Figure 4-53. Non-VSAM Wasted Space / Chart

Non-VSAM Wasted Space / Summary (STGECC)

Source: VCADAA file at the DETAIL timespan

Function: This report helps to identify departments or groups that are wasting DASD storage space due to poor allocation practices. These problems can be caused either by allocating more space than is necessary, or by using an inefficient block size. The Storage Administrator can analyze either or both of these causes with the same report. Correction of the problems detected by this report may recover a large amount of storage space with a small amount of effort.

Features: The left side of the report shows total space space wasted by each group due to the two causes described above, and also displays the amount of space wasted in megabytes by the most wasteful data set in the group. The middle portion of the report shows statistics related to poor block sizes, while the right side of the report summarizes waste due to overallocation of space. The owner of the data can be defined by one or more data set nodes, one or more account fields, or through an external table. This report can be run regularly to monitor wasted space, with more detailed reporting run when problems are detected.

Generation of the summary report is automatic.

SMS

Issues: Users planning for SMS could identify those groups with the worst space usage values, as they may be the best candidates for the immediate benefits of SMS. As groups migrate to SMS control, regular use of this report would verify that the expected space savings from SMS are being realized.

COLUMN DESCRIPTIONS

Data Owner

The value used as the summarization key for this report. Depending on the values specified on the report option screen, this will represent one or more VCA account elements, one or more nodes of the data set name, or will contain a value derived from your Data Set Group table.

Total Data Sets

The total number of data sets found for this summarization group in the data being analyzed. This value will be affected by the filtering criteria specified on the report option screen, especially if Max Data Sets was coded. The Max Data Sets option allows only a specified number of the worst data sets to be selected, in terms of wasted space. This option is valuable for reducing the size of the detail report, but will skew the results if you are trying to obtain an overview of all users.

Total MBytes Wasted

The total amount of DASD space being wasted by the data sets in this summarization group, expressed in millions of bytes (megabytes). This includes the space wasted due to poor block size (Blksize MBytes Wasted column), plus the space wasted due to over-allocation of space (OverAloc MBytes Wasted column).

Blksize MBytes Wasted

The total amount of DASD space being wasted by the data sets in this summarization group due to poor block size. For each data set, waste per track is first calculated by multiplying the block size by the blocks per track, and subtracting that value from the maximum track capacity for the device. That value is then multiplied by the number of tracks being used (not allocated) by the data set. This value will be zero if ALLOC analysis was specified on the report option screen when generating this report.

OverAloc MBytes Wasted

The total amount of DASD space being wasted by the data sets in this summarization group due to over-allocation of space. This is calculated for each data set by subtracting tracks used from tracks allocated, and multiplying that amount by the maximum track capacity for the device. This value will be zero if BLOCK analysis was specified on the report option screen when generating this report.

Worst Data Set

The amount of DASD space, expressed in millions of bytes (megabytes) being wasted by the worst data

set within this summarization group. This would be the data set that is wasting the largest amount of space of all data sets within the summarization group. You may want to run this report again using different values on the report option screen to identify the names of individual data sets that are wasting space.

Min Block Size

The smallest block size value found for all the data sets in this summarization group. This helps to identify groups that may be using small block sizes, which tend to waste space.

Avg Block Size

The average block size found for all the data sets in this summarization group. This is calculated as the sum of the block size value for each data set, divided by the number of data sets in the summarization group. This helps to identify groups that may be using small block sizes, which tend to waste space.

Avg# Blks Trk

The average number of blocks being stored on each DASD track for all the data sets in this summarization group. This is calculated as the sum of the number of blocks per track for each data set, divided by the number of data sets in the summarization group. This helps to identify groups that may be using poor block size values, as they will tend to write a larger number of blocks per track.

Avg % Track Waste

The average percentage of each track being wasted by the data sets in this summarization group because of block size factors. This is calculated by computing the maximum and actual number of bytes used per track for each data set, summing both of these values for all data sets in the summarization group, and then dividing the difference between the two sums by the sum of maximum space. This helps to identify groups that may be using poor block size values, as they will have a large percentage value in this column.

Total Tracks Alloc

The total number of tracks being allocated by all

the data sets in this summarization group.

Total Tracks Used

The total number of tracks actually being used to contain data for each of the data sets in this summarization group. The difference between the Total Tracks Alloc value and this value indicate the number of tracks that are being wasted by data sets in this group due to over-allocation.

Avg % Alloc Waste

The average percentage of the allocated tracks for all data sets within this summarization group that are not being used to contain data. This is calculated by subtracting the Total Tracks Used value from the Total Tracks Alloc value and dividing that by Total Tracks Alloc. This helps to identify groups that are over-allocating their data sets, as they will have large percentage values in this column.

NON-VSAM WASTED SPACE (STGECC)													4
CA MICS StorageMate - SUMMARY													
ABC CORPORATION													
----- SYSID=SYS7 -----													
Data Owner	Total Data Sets	Total MBytes Wasted	Blksize MBytes Wasted	OverAlloc MBytes Wasted	Worst Data Set	Min Block Size	Avg Block Size	Avg# Blks Trk	Avg % Track Waste	Total Tracks Alloc	Total Tracks Used	Avg % Alloc Waste	
CUSTOMER-SUPPORT	4	271	121	150	141	0	32760	1	30	11760	8560	27	
DEVELOPMENT	20	826	316	510	114	6160	31430	1	41	33046	22154	33	
DISTRIBUTION-LIBS	3	168	0	168	117	6160	15027	5	91	3600	0	100	
MARKETING-SALES	61	2971	990	1981	290	0	26774	6	54	91649	49448	46	
SYSTEM-DATA	12	1346	192	1154	620	0	12191	15	93	33935	9363	72	

SYSID	100	5582	1619	3963						173990	89525		
	=====	=====	=====	=====						=====	=====		
	100	5582	1619	3963						173990	89525		

Figure 4-54. Non-VSAM Wasted Space / Summary

Non-VSAM Wasted Space / Detail (STGECC)

Source: VCADAA file at the DETAIL timespan

Function: This report helps to identify users or departments that are wasting DASD storage space due to poor allocation practices. These problems are caused either by allocating more space than is necessary, or by using an inefficient block size. The Storage Administrator can analyze either or both of these causes with the same report. Correction of the problems detected by this report may recover a large amount of storage space with a small amount of effort.

Features: The middle portion of the report provides a detailed analysis of the effectiveness of the block size being used for each data set, including the number of blocks that will fit on one track, the percentage of each track that will be wasted, plus a suggested better value. The right side of the report shows waste due to overallocation of space. The user can limit the number of data sets, so that only the most wasteful data sets will appear. The report can also be sorted by space wasted or by data set name.

SMS

Issues: Users planning to use the system-determined block size option may be able to simulate the effects of more efficient block sizes, and thus help justify SMS and DFP. As DASD space may be at a premium during SMS implementation, this report could help ensure that critical free space is not wasted by poor allocations.

COLUMN DESCRIPTIONS

Data Set Name

The name of the data set. Depending on the values coded on the report option screen, the sequence of this report will be by Data Set Name or by the Total MByte Waste column.

Volume Serial Number

The volume serial number of the storage medium.

Total MByte Waste

The total amount of DASD space being wasted by

this data set, expressed in millions of bytes (megabytes). This includes the space wasted due to poor block size (Block MByte Waste column), plus the space wasted due to over-allocation of space (Alloc MByte Waste column). Depending on the values coded on the report option screen, the sequence of this report will be by Data Set Name or by the Total MByte Waste column.

Block Size

The maximum length of a physical block written into this data set.

LRECL

The maximum length of each logical record in this data set.

RECFM

Specifies the characteristics of the records in the data set. The records can be fixed length (F), variable length (V), or undefined length (U).

Blocks Per Track

The number of blocks of this data set that will fit on one DASD track. Multiplying this number by the Block Size will result in a determination of the number of bytes actually used on each track, and serves as the basis for determining waste due to poor Block Size. If this value is zero, the Blocks Per Track could not be calculated either because of a zero Block Size, or because the TRKCALC service ended in error. If the Block Size is not zero, examine the MICSLOG for any error messages related to TRKCALC.

% Each Track Unused

Based on the current Block Size, the percentage of each DASD track that is not being used. The larger this value, the more space that is being wasted. If this value is 100%, the value could not be calculated either because of a zero Block Size, or because the TRKCALC service ended in error. If the Block Size is not zero, examine the MICSLOG for any error messages related to TRKCALC.

Sugg. Block Size

A suggested Block Size for this data set that would cause more efficient use of each track. If this value is blank, the suggested size is either

already being used, or it could not be calculated due to the characteristics of the data set.

Block MByte Waste

The Non-VSAM Wasted Space report identifies wasted space caused by poor block sizes, overallocation of space, or both. If only waste by overallocation was requested, this column will contain zeros.

Tracks Alloc

The number of tracks allocated to this data set on the current volume.

Tracks Used

The number of tracks actually used by this data set on the current volume.

Alloc MByte Waste

The Non-VSAM Wasted Space report identifies wasted space caused by poor block sizes, overallocation of space, or both. If only waste by poor block size was requested, this column will contain zeros.

NON-VSAM WASTED SPACE (STGECC)												1
CA MICS StorageMate - DETAIL												
ABC CORPORATION												
----- SYSID=SYS7 -----												
Data Set Name	Volume Serial Number	Total MByte Waste	Block Size	LRECL	RECFM	Blocks Per Track	Sugg. Block Size	% Each Track Unused	Block MByte Waste	Tracks Alloc	Tracks Used	Alloc MByte Waste
SYS7.HASPACE	SP00L1	620	0	0	V	0	23476	100	0	13200	0	620
MVSRGS.IPL4BB.D900413	SYS080	290	0	0	U	0		100	289	6165	6152	1
MVSEGP.AKD000.STAN	PAG406	188	0	0	U	0		100	188	4005	3998	0
ITS035.SNA.MICS.DETAIL	ITS003	141	32760	32756	U	1		31	36	4800	2569	105
QCAM.ABC9004U.DETAIL	SPAR01	140	32760	32756	U	1		31	0	3000	10	140
AKD000.FUT8910.TEST.LIST	SPAR04	117	6160	80	FB	7	23440	9	0	2500	0	117
AKD052.MAMICSUP.DETAIL	LNR934	114	32760	32756	U	1		31	66	5700	4668	48
QCAM.TMP9004D.DETAIL	SPAR05	112	32760	32756	U	1		31	15	3130	1062	97
SYS7.TARGETMC	PAG405	111	200	200	FB	68	23400	72	76	3000	2264	35
SYS7.TMC	MCAT04	107	200	200	F	68	23400	72	83	3000	2485	24
QCAM.ALL.DOC.TEXT	SPAR02	42	6160	137	VB	7	23476	9	0	900	14	42
OPRSEH.SMF.DETAIL	LNR927	41	32760	32756	U	1		31	0	900	23	41
AKD052.MAMICSUP.WEEKS	LNR921	41	32760	32756	U	1		31	23	1983	1596	18
ONLSMS.P8910IMS.WEEKS	LNR928	40	32760	32756	U	1		31	0	850	2	40
CAPMGG.DETAIL	LNR804	39	32760	32756	U	1		31	10	1350	739	29
PASSRB.PAS4.INT.BACKUP.DATABASE	DAT002	39	32760	32756	U	1		31	25	2055	1755	14
AKD052.MAMICSUP.MONTHS	LNR919	39	32760	32756	U	1		31	38	2700	2680	1
AKD182.SAMES.SPACE	SPAR01	106	6160	80	FB	7	23440	9	0	2250	0	106
QCAM.NAS9004U.DETAIL	LNR918	105	32760	32756	U	1		31	3	2400	223	102
SYS7.DUMP01	SMF411	95	4104	4104	F	10	20520	13	0	2025	0	95
SYS7.DUMP02	SMF411	95	4104	4104	F	10	20520	13	0	2025	0	95
SYS7.DUMP03	SMF411	95	4104	4104	F	10	20520	13	0	2025	0	95
ONLSSE.IMS.DLOG.V310.FEDX.S0C4@0C	SPAR02	90	32760	32756	VB	1		31	90	6345	6345	0
PASSRH.PAS4.OLD.INFO.HISTORY	DAT001	78	32760	32756	U	1		31	45	3900	3201	33
OPRSES.SRL.DETAIL	LNR922	77	32760	32756	U	1		31	0	1650	3	77
CSH.PRODUCTS.SASDB	LNR801	77	32760	32756	U	1		31	36	3375	2505	41
QCAM.USA9004D.DETAIL	SPAR03	70	32760	32756	U	1		31	2	1600	147	68
AKD052.MAMICSUP.DAYS2	LNR806	67	32760	32756	U	1		31	2	1500	106	65
AKD052.MAMICSUA.DETAIL	LNR919	65	32760	32756	U	1		31	28	2750	1953	37
AKD052.MAMICSUP.DAYS	LNR933	63	32760	32756	U	1		31	53	3930	3722	10
QCAM.INT9004F.DETAIL	LNR934	61	32760	32756	U	1		31	4	1500	277	57
ONLSSE.P8910IMS.DETAIL	LNR922	61	32760	32756	U	1		31	0	1300	2	61
MSSRRP.CP.DETAIL	EIS001	60	32760	32756	U	1		31	21	2325	1500	39
SYS7.LINKLIB	SP40L1	59	32760	0	U	1		31	10	1755	721	49
QCAM.STST8910D.DETAIL	SPAR03	58	32760	32756	U	1		31	5	1500	367	53

SYSID		5582							1619	173990	89525	3963
		=====							=====	=====	=====	=====
		5582							1619	173990	89525	3963

Figure 4-55. Non-VSAM Wasted Space / Detail

SMS Class Usage / Chart (STGELE)

Source: VCADAA file at the DETAIL timespan
 VCA_VS file at the DETAIL timespan
 HSMMIG file at the DETAIL timespan
 HSBAC file at the DETAIL timespan

Function: Graphically represents the usage of DFSMS classes by department.

Monitors the use of various DFSMS classes. It provides a view of what departments are using a particular class, detects potential violations, and ensures that DFSMS is assigning classes correctly.

DFSMS classes are set either by the user or by the DFSMS ACS routines, and are used to group data sets with similar attributes and needs.

Features: The user specifies the type of class to be summarized (data class, storage class, storage group, or management class). The chart shows the class names and the amount of space used by each department.

Users can report class usage based on one or more of the account fields, one or more data set name nodes, or by using a user-defined table that translates the data set name prefix into an owner description.

Both color graphics and printer charts are supported.

SMS Issues: This report is useful during SMS implementation to ensure that your SMS ACS routines are assigning class names as designed. After SMS implementation, the report is useful for monitoring the usage of SMS classes by various departments/users.

COLUMN DESCRIPTIONS

Storage Class

The name of the DFSMS Storage Class for which this chart applies. A value of '(NONE)' is used to represent data sets that have not been assigned to

any class. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this chart.

Management Class

The name of the DFSMS Management Class for which this chart applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any class. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this chart.

Data Class

The name of the DFSMS Data Class for which this chart applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any class. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this chart.

Storage Group

The name of the DFSMS Storage Group for which this chart applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any group. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this chart. Because Storage Group is not available for data sets managed by HSM, values of 'HSM LV1' (migration level 1), 'HSM LV2' (migration level 2) and 'HSM BACK' (HSM backup) will appear in this column for HSM data.

Pie Slices

Each slice of the pie represents the total amount of space allocated by the users of this DFSMS class or group, expressed in millions of bytes (megabytes).

Pie Labeling

The label for each pie slice contains the name of the group represented by that slice. These summarization groups can be constructed using one or more VCA account elements, one or more nodes of the data set name, or using a value derived from your Data Set Group table.




Figure 4-56. SMS Class Usage / Chart

SMS Class Usage / Summary (STGELE)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan

Function: Summarizes the installation's usage of various DFSMS classes.

Helps the user track growth rates of SMS classes, detect potential problems, and ensure that DFSMS is assigning classes correctly.

DFSMS classes are set either by the user or by the DFSMS ACS routines, and are used to group data sets with similar attributes and needs.

Features: The user specifies the type of class to be summarized: (data class, storage class, storage group, or management class). The report shows the names of the various classes found, the number of data sets using each class, and the amount of space this represents, both in terms of space allocated and space used. Flexible filtering criteria, plus summarization at both the class and system levels, allows many different views of SMS class usage information.

The ability to include data from real DASD plus HSM allows the ability to monitor SMS classes as data sets migrate through various storage levels.

SMS Issues: This report is useful during SMS implementation to ensure that your SMS ACS routines are assigning class names as designed. After SMS implementation, the report is useful for monitoring the usage of SMS classes in various environments.

COLUMN DESCRIPTIONS

Storage Class

The name of the DFSMS Storage Class for which this summarization applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any class. The report options screen is used to select which of the four DFSMS

constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this summary report.

Management Class

The name of the DFSMS Management Class for which this summarization applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any class. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this summary report.

Data Class

The name of the DFSMS Data Class for which this summarization applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any class. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this summary report.

Storage Group

The name of the DFSMS Storage Group for which this summarization applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any group. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this summary report. Because Storage Group is not available for data sets managed by HSM, values of 'HSM LV1' (migration level 1), 'HSM LV2' (migration level 2) and 'HSM BACK' (HSM backup) will appear in this column for HSM data.

Total Data Sets

The total number of data sets found in the input data that were assigned to this class or group.

Non-VSAM Data Sets

The total number of data sets found in the input data that were assigned to this class or group and were defined as having a data set organization other than VSAM. This value includes only data sets that still reside on primary DASD (no HSM data).

VSAM Data Sets

The total number of data sets found in the input data that were assigned to this class or group and were defined as having a VSAM organization. This value includes only data sets that still reside on primary DASD (no HSM data).

HMIG Data Sets

The total number of data sets found in the input data that were assigned to this class or group and had been migrated to HSM level 1 or level 2 volumes.

HBAC Data Sets

The total number of data sets found in the input data that were assigned to this class or group and belonged to data sets residing on the HSM backup volumes. Each unique data set is only included once in this count, even if multiple backup copies exist.

Total Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group, expressed in millions of bytes (megabytes).

Non-VSAM Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group and belonged to data sets with a data set organization other than VSAM. This value only includes space allocated by data sets still residing on primary DASD (no HSM data).

VSAM Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group and belonged to data sets with a VSAM data set organization. This value is expressed in millions of bytes (megabytes), and only represents data sets residing on primary DASD (no HSM data).

H-Migrate Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group and was allocated by data sets that had been migrated to HSM level 1 or level 2 volumes.

H-Backup Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group and was allocated to HSM backup storage. This value may include multiple backups of the same data set.

SMS CLASS USAGE (STGELE)											
CA MICS StorageMate - SUMMARY											
ABC CORPORATION											
----- System Identifier=SYSA -----											
Storage Class	Total Data Sets	Non-VSAM Data Sets	VSAM Data Sets	HMIG Data Sets	HBAC Data Sets		Total Megabytes Allocated	Non-VSAM Megabytes Allocated	VSAM Megabytes Allocated	H-Migrate Megabytes Allocated	H-Backup Megabytes Allocated
(NONE)	12461	8218	1732	2486	25		38430.2	28320.9	8288.4	1778.9	41.9
CRITICAL	1732	485	406	294	547		7099.7	2012.1	2376.2	352.3	2359.1
STANDARD	43056	12283	1863	18021	10889		102313.8	52869.4	3771.7	18807.2	26865.5
TECH1	3	1	0	2	0		0.1	0.1	0.0	0.0	0.0
TECH3	4	2	0	1	1		5.7	1.8	0.0	2.8	1.1
-----	-----	-----	-----	-----	-----		-----	-----	-----	-----	-----
SYSID	57256	20989	4001	20804	11462		147849.4	83204.3	14436.2	20941.2	29267.6
	=====	=====	=====	=====	=====		=====	=====	=====	=====	=====

Figure 4-57. SMS Class Usage / Summary

SMS Class Usage / Owner (STGELE)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan

Function: This report allows the user to summarize the usage of various DFSMS classes by various departments and users.

The report shows the user which departments are using a particular class and helps to detect potential violations, and to make sure DFSMS is assigning classes correctly.

DFSMS classes are set either by the user or by the DFSMS ACS routines, and are used to group data sets with similar attributes and needs.

Features: The user specifies the class to be summarized: storage group, data class, storage class, or management class. The report shows the names of the various users of each class, the number of data sets using the class, and the amount of space this represents, both in terms of space allocated and space used.

Users can report class usage based on one or more of the account fields, one or more data set name nodes, or by using a user-defined table that translates the data set name prefix into an owner description.

The ability to include data from real DASD plus HSM allows the ability to monitor SMS classes as data sets migrate through various storage levels.

SMS Issues: This report is useful during SMS implementation to ensure that your SMS ACS routines are assigning class names as designed. After SMS implementation, the report is useful for monitoring the usage of SMS classes by various departments/users.

COLUMN DESCRIPTIONS

Storage Class

The name of the DFSMS Storage Class for which this summarization applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any class. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this summary report.

Management Class

The name of the DFSMS Management Class for which this summarization applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any class. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this summary report.

Data Class

The name of the DFSMS Data Class for which this summarization applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any class. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this summary report.

Storage Group

The name of the DFSMS Storage Group for which this summarization applies. A value of '(NONE)' is used to represent data sets that have not been assigned to any group. The report options screen is used to select which of the four DFSMS constructs (Storage Class, Management Class, Data Class, or Storage Group) will appear on this summary report. Because Storage Group is not available for data sets managed by HSM, values of 'HSM LV1' (migration level 1), 'HSM LV2' (migration level 2) and 'HSM BACK' (HSM backup) will appear in this column for HSM data.

Data Owner

The value used as the summarization key for this report. Depending on the values specified on the report option screen, this will represent one or more VCA account elements, one or more nodes of

the data set name, or will contain a value derived from your Data Set Group table.

Total Data Sets

The total number of data sets found in the input data that were assigned to this class or group.

Non-VSAM Data Sets

The total number of data sets found in the input data that were assigned to this class or group and were defined as having a data set organization other than VSAM. This value includes only data sets that still reside on primary DASD (no HSM data).

VSAM Data Sets

The total number of data sets found in the input data that were assigned to this class or group and were defined as having a VSAM organization. This value includes only data sets that still reside on primary DASD (no HSM data).

HMIG Data Sets

The total number of data sets found in the input data that were assigned to this class or group and had been migrated to HSM level 1 or level 2 volumes.

HBAC Data Sets

The total number of data sets found in the input data that were assigned to this class or group and belonged to data sets residing on the HSM backup volumes. Each unique data set is only included once in this count, even if multiple backup copies exist.

Total Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group, expressed in millions of bytes (megabytes).

Non-VSAM Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group and belonged to data sets with a data set organization other than VSAM. This value only includes space allocated by data sets still residing on primary DASD (no HSM data).

VSAM Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group and belonged to data sets with a VSAM data set organization. This value is expressed in millions of bytes (megabytes), and only represents data sets residing on primary DASD (no HSM data).

H-Migrate Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group and was allocated by data sets that had been migrated to HSM level 1 or level 2 volumes.

H-Backup Megabytes Allocated

The total amount of allocated space found in the input data that was assigned to this class or group and was allocated to HSM backup storage. This value may include multiple backups of the same data set.

SMS CLASS USAGE (STGELE)										
CA MICS StorageMate - OWNER										
ABC CORPORATION										
----- System Identifier=SYSA Storage Class=(NONE) -----										
Data Owner	Total Data Sets	Non-VSAM Data Sets	VSAM Data Sets	HMIG Data Sets	HBAC Data Sets	Total Megabytes Allocated	Non-VSAM Megabytes Allocated	VSAM Megabytes Allocated	H-Migrate Megabytes Allocated	H-Backup Megabytes Allocated
FINANCE	6685	4210	566	1907	2	8715.2	6789.5	1160.2	764.1	1.4
SUPPORT	2397	2081	272	44	0	17138.8	15736.2	1255.4	147.2	0.0
SALES	1016	353	423	238	2	1255.5	680.6	323.6	245.0	6.2
CORPORATE	870	540	330	0	0	2534.6	1036.2	1498.4	0.0	0.0
SYS-PROGS	864	576	75	213	0	3617.7	3002.7	101.9	513.1	0.0
DATA-BASE-PROD	566	406	55	84	21	5085.1	1006.8	3934.5	109.4	34.4
MARKETING	42	37	5	0	0	65.8	52.6	13.3	0.0	0.0
DATA-BASE-TEST	21	15	6	0	0	17.4	16.3	1.1	0.0	0.0
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
STORSTGC	12461	8218	1732	2486	25	38430.2	28320.9	8288.4	1778.9	41.9
----- System Identifier=SYSA Storage Class=CRITICAL -----										
Data Owner	Total Data Sets	Non-VSAM Data Sets	VSAM Data Sets	HMIG Data Sets	HBAC Data Sets	Total Megabytes Allocated	Non-VSAM Megabytes Allocated	VSAM Megabytes Allocated	H-Migrate Megabytes Allocated	H-Backup Megabytes Allocated
SUPPORT	706	206	144	161	195	2163.2	928.9	603.4	237.6	393.2
FINANCE	325	93	55	78	99	366.5	113.3	86.5	84.2	82.5
DATA-BASE-PROD	285	111	24	21	129	1029.7	576.8	88.9	24.7	339.4
SALES	222	54	83	7	78	3368.0	312.1	1560.9	0.9	1494.0
SYS-PROGS	178	19	94	26	39	156.2	80.4	24.3	4.8	46.7
CORPORATE	10	0	6	1	3	12.9	0.0	12.1	0.1	0.7
MARKETING	4	2	0	0	2	0.7	0.6	0.0	0.0	0.1
DATA-BASE-TEST	2	0	0	0	2	2.4	0.0	0.0	0.0	2.4
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
STORSTGC	1732	485	406	294	547	7099.7	2012.1	2376.2	352.3	2359.1

Figure 4-58. SMS Class Usage / Owner

SMS Class Usage / Detail (STGELE)

Source: VCADAA file at the DETAIL timespan
 VCA_VS file at the DETAIL timespan
 HSMIMG file at the DETAIL timespan
 HSBAC file at the DETAIL timespan

Function: Provides a detailed listing of an installation's data sets, and the names of the SMS constructs (storage class, management class, data class, and storage group) assigned to them.

DFSMS classes are set either by the user or by the DFSMS ACS routines, and are used to group data sets with similar attributes and needs.

This report is useful in conjunction with the other SMS class usage reports, to obtain more detailed information about a particular class. For example, if the owner report showed that a particular department were using an SMS class they were not authorized for, this report could be used to list all the individual data sets assigned to that class.

Features: The user specifies the type of class to be summarized: (data class, storage class, storage group, or management class). Flexible filtering criteria allows for many views of the data, from very broad to very narrow.

The ability to include data from real DASD plus HSM allows the ability to monitor SMS classes as data sets migrate through various storage levels.

SMS Issues: This report is useful during SMS implementation to ensure that your SMS ACS routines are assigning class names as designed, and to investigate potential problems in this area. After SMS implementation, the report is useful as part of the ongoing monitoring and tuning of the SMS system.

COLUMN DESCRIPTIONS

Storage Class

The name of the DFSMS Storage Class assigned to

this data set. A value of '(NONE)' indicates that no value has been assigned. Depending on the values coded on the report options screen, one construct name (Storage Class, Management Class, Data Class, and Storage Group) will be selected as the summarization element, while the other three will appear in the body of the report.

Management Class

The name of the DFSMS Management Class assigned to this data set. A value of '(NONE)' indicates that no value has been assigned. Depending on the values coded on the report options screen, one construct name (Storage Class, Management Class, Data Class, and Storage Group) will be selected as the summarization element, while the other three will appear in the body of the report.

Data Class

The name of the DFSMS Data Class assigned to this data set. A value of '(NONE)' indicates that no value has been assigned. Depending on the values coded on the report options screen, one construct name (Storage Class, Management Class, Data Class, and Storage Group) will be selected as the summarization element, while the other three will appear in the body of the report.

Storage Group

The name of the DFSMS Storage Group assigned to this data set. A value of '(NONE)' indicates that no value has been assigned. Depending on the values coded on the report options screen, one construct name (Storage Class, Management Class, Data Class, and Storage Group) will be selected as the summarization element, while the other three will appear in the body of the report. Because Storage Group is not available for data sets managed by HSM, values of 'HSM LV1' (migration level 1), 'HSM LV2' (migration level 2) and 'HSM BACK' (HSM backup) will appear in this column for HSM data.

Data Set Name

The name of the data set.

Volume Serial Number

The volume serial number of the storage medium. For data sets residing on HSM storage, values of

HSMLV1 (migration level 1), HSMLV2 (migration level 2), and HSBAC (backup) will be used.

DSORG

The organization of the data set.

Space Alloc (KB)

The number of kilobytes allocated to this data set on the current volume.

Space Used (KB)

The number of kilobytes used by this data set on the current volume. Space used will always equal space allocated for data sets that reside on HSM storage, unless the PRIMARY option was specified for HSM space reporting.

Create Date

The date that the data set was first allocated to a DASD volume. This value will always be missing (.) for data sets residing on HSM backup storage.

Date Last Used

The date that the data set was last referenced. This value will be missing (.) for data sets residing on HSM backup storage.

SMS CLASS USAGE (STGELE)									
CA MICS StorageMate - DETAIL									
ABC CORPORATION									
----- System Identifier=SYS7 Management Class=CRITICAL -----									
Data Set Name	Volume Serial Number	DSORG	Space Alloc (KB)	Space Used (KB)	Create Date	Date Last Used	Storage Class	Data Class	Storage Group
AJA003.ISPPROF	LNR806	PO	95	47	03APR90	.	PROD01		PRODWRK1
BMANAGE.NDVRC1.PROD.CONLIB	LNR800	PO	5697	4558	22MAY90	21AUG90	PROD01		PRODWRK1
BMANAGE.NDVRC1.PROD.PROCLIB	LNR803	PO	95	95	22MAY90	21AUG90	PROD01		PRODWRK1
BMANAGE.NDVRC1.TEST.ZAPS	LNR800	PO	285	237	01JUN90	07AUG90	PROD01		PRODWRK1
BMANAGE.NDVVMVS.TEST.SOURCE	LNR803	PO	142	95	31JUL90	16AUG90	PROD01		PRODWRK1
BMANAGE.NDVR35.CLIST	LNR806	PO	95	95	22MAY90	06AUG90	PROD01		PRODWRK1
PROCESS.A.CHECKPT.DATA	LNR803	PS	237	95	21AUG90	21AUG90	PROD01		PRODWRK1
PROCESS.A.CNTL	LNR803	PO	570	475	20AUG90	21AUG90	PROD01		PRODWRK1
PROCESS.A.RESTART.CNTL	LNR803	***	237	0	21AUG90	.	PROD01		PRODWRK1
PROCESS.A.YEARS	LNR803	DA	95	95	21AUG90	21AUG90	PROD01		PRODWRK1
PROCESS.C.CLIST	LNR804	PO	1614	950	14OCT88	21AUG90	PROD01		PRODWRK1
PROCESS.C.CNTL	LNR804	PO	1424	950	26JAN89	20AUG90	PROD01		PRODWRK1
PROCESS.C.INCLLIB	LNR804	PO	1424	1044	14JUN90	21AUG90	PROD01		PRODWRK1
PROCESS.C.ISPMLIB	LNR804	PO	1709	1377	17OCT88	21AUG90	PROD01		PRODWRK1
PROCESS.C.ISPPLIB	LNR803	PO	17801	15097	17AUG90	21AUG90	PROD01		PRODWRK1
PROCESS.C.PARMS	LNR804	PO	1044	380	17OCT88	21AUG90	PROD01		PRODWRK1
PROCESS.C.PROCLIB	LNR804	PO	1044	712	17OCT88	21AUG90	PROD01		PRODWRK1
PROCESS.C.PROTOLIB	LNR804	PO	2136	1757	14JUN90	21AUG90	PROD01		PRODWRK1
PROCESS.C.PSP.MACLIB	LNR803	PO	47	47	24JUL90	17AUG90	PROD01		PRODWRK1
PROCESS.COMPLEX.CHANGES	LNR803	PO	47	47	03MAY90	26JUN90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4282.PSP.SOURCE	LNR803	PO	47	47	17MAY90	17JUL90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4283.PSP.DOC.TEXT	LNR800	PO	47	47	17JUL90	20JUL90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4290.PSP.ASM	LNR803	PO	237	95	18JUN90	21AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4290.PSP.CLIST	LNR803	PO	47	47	06JUL90	21AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4290.PSP.GENLIB	LNR801	PO	47	47	06JUL90	21AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4290.PSP.ISPMLIB	LNR803	PO	95	95	30APR90	21AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4290.PSP.ISPPLIB	LNR803	PO	237	237	30APR90	09AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4290.PSP.ISPTLIB	LNR803	PO	285	285	30APR90	21AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4290.PSP.PC.TEXT	LNR803	PO	237	190	30APR90	21AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.DRK4290.PSP.USOURCE	LNR806	PO	47	47	06JUL90	21AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.RAK4223.PSP.DOC.TEXT	LNR804	PO	47	47	11APR90	11JUN90	PROD01		PRODWRK1
PROCESS.COMPOSIT.RAK4230.PSP.ASM	LNR800	PO	95	95	05JUL90	21AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.RAK4230.PSP.FIXES	LNR803	PO	47	47	05JUL90	07AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.RAK4230.PSP.ISPPLIB	LNR807	PO	47	47	05JUL90	20AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.RAK4230.PSP.MACAUTOS	LNR801	PO	47	47	05JUL90	21AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.SW04254.PSP.DIC.TEXT	LNR804	PO	47	47	11JUL90	10AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.SW04254.PSP.LOAD	LNR803	PO	47	47	11JUL90	14AUG90	PROD01		PRODWRK1
PROCESS.COMPOSIT.SW04254.PSP.USOURCE	LNR807	PO	47	47	11JUL90	10AUG90	PROD01		PRODWRK1
PROCESS.D.CHECKPT.DATA	LNR803	PS	237	95	21AUG90	21AUG90	PROD01		PRODWRK1
PROCESS.D.CNTL	LNR803	PO	570	475	20AUG90	21AUG90	PROD01		PRODWRK1
PROCESS.D.WEEKS	LNR803	DA	95	95	21AUG90	21AUG90	PROD01		PRODWRK1
PROCESS.DDRK4280.NWI.MONTHS	LNR807	DA	2516	2516	27NOV89	05APR90	PROD01		PRODWRK1

Figure 4-59. SMS Class Usage / Detail

SMS Class Usage / Descriptive Statistics (STGELE)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan

Function: Provide simple descriptive statistics for both Age and Allocated Space (data set size) for each cluster. This helps determine the general size and shape of individual clusters, and can additionally help identify the presence of outliers.

Features: This report consists of what is commonly called a "five number summary" of both Age and Data Set Size as opposed to the often used mean and standard deviation measures. This approach was chosen because it is less affected by non-normal statistical distributions of the input data values and is typically used in engineering applications for that very reason. Aside from the descriptive statistics, the total number of data sets and allocated space is given to quantify the affect of the cluster with respect to the total population, and a set of weighted allocation measurements is provided as an alternative metric.

SMS N/A

Issues:

COLUMN DESCRIPTIONS

Cluster

The number of the cluster, from one to n, assigned by the clustering application.

Number of Datasets

The number of data sets within the cluster.

Age (in days) Statistical Summary

The Age is the time in days since the last reference of the data set being processed. It is computed by subtracting the last reference date from the current date. If the result is missing, then the same calculation is attempted using the data set creation date. Observations that still have missing values after both attempts are deleted. The Age values are displayed as Minimum, 1st Quartile (25th pctl.), median, 3rd Quartile (75th pctl.) and Maximum.

Allocated Space (size) in Megabytes

This measurement is derived from element DAAALLSP and presented as Minimum, 1st Quartile (25th pctl.), median, 3rd Quartile (75th pctl.), Maximum, and total values. This total size depicts how much allocated space is represented by the cluster.

Weighted Allocated Space

The allocated space of each data set is first weighted by its calculated Age and then accumulated together within each cluster. The accumulated measurements for each cluster are then presented both as an Average (mean) and Total values. This is done to quantify the effect of the size and ages of all data sets taken together and provides a simple method to understand where the "pain" is. The expression "pain" refers to the potential impact one set of data sets might have on the installation as a whole. For example, a cluster representing a group of poorly managed data sets might consist of many small data sets that are fairly new, and therefore cause minimal "pain." Conversely, a group of rather large data sets that are old can easily result in a large impact on the installation.

SMS CLASS USAGE (STGELE)															4
CA MICS StorageMate															
AFTER COMPLEX SAS V8.2															
Descriptive Statistics for Age & Dataset Size															
Report Date: 06/25/2004															
Cluster	No. of Datasets	Age (Days)					Allocated Space (MB)					Total	Wgtd. Alloc. Space (MB)		
		Minimum	1st Qrtl.	Median	3rd Qrtl.	Maximum	Minimum	1st Qrtl.	Median	3rd Qrtl.	Maximum		Average	Total	
1	29,480	20	23	35	97	372	0	0	1	2	16	48,349	104	3,074,114	
2	9,258	289	401	484	617	865	0	0	0	1	17	7,283	412	3,813,085	
3	3,150	863	997	1,172	1,380	1,903	0	0	0	1	22	3,576	1,346	4,241,299	
4	1,656	20	22	31	107	2,052	16	20	27	41	61	50,573	3,945	6,532,345	
5	500	21	23	31	90	2,030	62	76	85	108	156	46,633	10,554	5,276,859	
6	698	1,910	2,145	2,486	2,796	6,935	0	0	0	1	28	534	1,978	1,380,403	
7	205	21	22	31	56	2,380	160	172	212	255	325	45,483	20,271	4,155,606	
8	104	21	21	28	37	967	330	357	425	486	640	44,878	36,830	3,830,367	
9	4	21	22	26	48	66	2,628	2,716	2,963	3,168	3,216	11,770	99,406	397,622	
10	29	21	22	30	253	1,127	680	765	827	907	1,700	25,714	214,216	6,212,277	

Figure 4-59A SMS Class Usage / Descriptive Statistics

SMS Class Usage / Distribution of Allocated Space (STGELE)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan

Function: Provides a summary of the allocated space within each DFSMS Storage Group within each primary summarization group. The total number of data sets and allocated space for each summarization group is presented, and the amount of space distributed within each Storage Group is also displayed. You can quickly determine the overall "health" of the storage subsystem, by reviewing the mapping of allocations for each summarization group across the respective Storage Groups. The default value for the primary summarization group is Storage Class. Generally, each Storage Class should map to only a small number of groups, ideally one to one. When a given storage class maps to many groups, the definitions within the Storage Class should be reviewed and adjusted as required.

The report format is controlled by limiting the number of Storage Groups to be displayed in the report. This filtering is done by selecting those storage groups that represent at least 2 percent of all data sets and combining all others into a single definition. Additionally, all non-SMS managed data sets are grouped together as well.

Features: An additional element for summarizing and reporting the allocation data can be specified. For example, the DFSMS Management Class can be specified so that all allocations are mapped by management class within storage class to better verify the DFSMS definitions.

SMS Issues: After SMS implementation, the report is useful as part of the ongoing monitoring and tuning of the SMS system.

COLUMN DESCRIPTIONS

The valid data element choices for primary summarization key are: STORAGE (storage class); MANAGE (management class); DATA (data class); and ACCOUNT, one of the VCAACTx elements. The default value is STORAGE.

An additional element can be used for secondary summarization. The valid choices are: `MANAGE` (management class); `DATA` (data class); `ACCOUNT` (`VCAACTx`); and `NONE`. The default value is `NONE`, indicating that only primary summarization will be used.

Following is a list of field descriptions for the summarization elements.

Storage Class

The name of the DFSMS Storage Class assigned to this data set. A value of `'(NONE)'` indicates that no value has been assigned.

Management Class

The name of the DFSMS Management Class assigned to this data set. A value of `'(NONE)'` indicates that no value has been assigned.

Data Class

The name of the DFSMS Data Class assigned to this data set. A value of `'(NONE)'` indicates that no value has been assigned.

Account fields

These are the `VCAACTx` elements that were specified during VCA component installation. Any of the defined account codes may be chosen; however, the report design limits the display of any summarization elements to 12 bytes.

The DFSMS Storage Group, discussed below is not used for summarization purposes, but is used to display the distribution of allocated space.

Storage Group

The name of the DFSMS Storage Group assigned by the installation. Only those Storage Groups that represent at least 2 percent of the allocated space are selected and then only the top 8 groups are actually displayed. All others are grouped together as `"OTHER"`.

Number of Datasets

The number of data sets within the summarization group.

Total Space (MB)

The total allocated space, displayed in megabytes
within the summarization group.

SMS CLASS USAGE (STGELE)												1
CA MICS StorageMate												
AFTER COMPLEX SAS V8.2												
Distribution of Space within Storage Groups												
Report Date: 06/25/2004												
DFSMS Storage Groups												
DFSMS Storage Class Name	DFSMS Management Class Name	No. of Datasets	Total Space (MB)	(NONSMS)	DBDSALC	MEDDSALC	MISCSMS	SMLDSALC	SPCDSALC	TSOSTOR	UTILSTOR	(OTHER)
(NONE)	(NONE)	50	597	0	42	69	35	364	0	5	0	82
		50	597	0	42	69	35	364	0	5	0	82
(NONSMS)	(NONSMS)	25,615	113,933	113,933	0	0	0	0	0	0	0	0
		25,615	113,933	113,933	0	0	0	0	0	0	0	0
DBDSUSRS	DEFAULT	497	4,394	0	4,394	0	0	0	0	0	0	0
	DEV	35	577	0	577	0	0	0	0	0	0	0
	SPFPERM	27	15	0	15	0	0	0	0	0	0	0
	SPFTEMP	116	612	0	612	0	0	0	0	0	0	0
	STANDARD	205	2,398	0	2,398	0	0	0	0	0	0	0
	TSO	8	50	0	50	0	0	0	0	0	0	0
	WORK	3	66	0	66	0	0	0	0	0	0	0
		891	8,112	0	8,112	0	0	0	0	0	0	0
DEVUSRS	DEV	994	5,687	0	0	0	0	0	0	0	0	5,687
	WORK	13	5	0	0	0	0	0	0	0	0	5
		1,007	5,692	0	0	0	0	0	0	0	0	5,692
GENDSUSR	DEFAULT	204	3,180	0	0	0	0	0	0	0	0	3,180
	SPFPERM	14	4	0	0	0	0	0	0	0	0	4
	TSO	9	8	0	0	0	0	0	0	0	0	8
	WORK	1	1	0	0	0	0	0	0	0	0	1
		228	3,193	0	0	0	0	0	0	0	0	3,193
MEDDSUSR	(NONE)	23	56	0	0	56	0	0	0	0	0	0
	DEFAULT	1,126	24,554	0	0	24,554	0	0	0	0	0	0
	IPCSDUMP	3	37	0	0	37	0	0	0	0	0	0
	SPFPERM	24	13	0	0	13	0	0	0	0	0	0
	SPFTEMP	14	131	0	0	131	0	0	0	0	0	0
	SVCDUMPS	1	26	0	0	26	0	0	0	0	0	0
	TSO	25	124	0	0	124	0	0	0	0	0	0
	WORK	19	600	0	0	600	0	0	0	0	0	0
		1,235	25,541	0	0	25,541	0	0	0	0	0	0
(sample report truncated to one page for documentation purposes)												

45,084	284,794	113,933	8,155	25,609	13,386	17,010	5,888	58,684	31,180	10,949		

Figure 4-59B SMS Class Usage / Distribution of Allocated Space

SMS Class Usage / Analysis of Space Usage (STGELE)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan

Function: Provides a summary of the Age and Size of data sets and how each grouping (cluster) is actually mapped into each Storage Group. This determines how well the DFSMS storage definitions are actually being used and where potential problems might reside.

Simple statistics (minimum, average and maximum) for both Age and Size are provided to aid in this process.

The report format is controlled by limiting the number of storage groups to be displayed in the report. This filtering is done by selecting those storage groups that represent at least 2 percent of all data sets and combining all others into a single definition. Additionally, all non-SMS managed data sets are grouped together as well.

Features: An additional element for summarizing and reporting the allocation data can be specified. By default, the DFSMS Storage Class is used to summarize allocations across storage groups to better verify the DFSMS definitions.

SMS Issues: After SMS implementation, the report is useful as part of the ongoing monitoring and tuning of the SMS system.

COLUMN DESCRIPTIONS

The valid data element choices for primary summarization key are: STORAGE (storage class); MANAGE (management class); DATA (data class); and ACCOUNT, one of the VCAACTx elements. The default value is STORAGE.

Following is a list of field descriptions for the summarization elements.

Storage Class

The name of the DFSMS Storage Class assigned to this data set. A value of '(NONE)' indicates that no value has been assigned.

Management Class

The name of the DFSMS Management Class assigned to this data set. A value of '(NONE)' indicates that no value has been assigned.

Data Class

The name of the DFSMS Data Class assigned to this data set. A value of '(NONE)' indicates that no value has been assigned.

Account fields

These are the VCAACTx elements that were specified during VCA component installation. Any of the defined account codes may be chosen; however, the report design limits the display of any summarization elements to 12 bytes.

Storage Group

The name of the DFSMS Storage Group assigned by the installation. Only those Storage Groups that represent at least 2 percent of the allocated space are selected and then only the top 8 groups are actually displayed. All others are grouped together as "OTHER".

Number of Datasets

The number of data sets within the summarization group.

Age (in days) Statistical Summary Age is the number of days since the last reference of the data set being processed. It is computed by subtracting the last reference date from the current date. If the result is missing, then the same calculation is attempted using the data set creation date. Observations that still have missing values after both attempts are deleted. The Age values are displayed as Minimum, 1st Quartile (25th pctl.), median, 3rd Quartile (75th pctl.) and Maximum.

Allocated Space (size) in Megabytes This is derived from element DAAALLSP, presented as Minimum, Average (mean) and Maximum values, and additionally summarized. This total size depicts how much allocated space is represented by the cluster within a given storage group.

Weighted Allocated Space

The allocated space of each data set is first weighted by its calculated Age and then accumulated together within each cluster. The accumulated measurements for each cluster are then presented both as an Average (mean) and Total values. This is done to quantify the effect of the size and ages of all data sets taken together and provides a simple method to understand where the "pain" is. The expression "pain" refers to the potential impact one set of data sets might have on the installation as a whole. For example, a cluster representing a group of poorly managed data sets might consist of many small data sets that are fairly new, and therefore cause minimal "pain." Conversely, a group of rather large data sets that are old can easily result in a large impact on the installation.

SMS CLASS USAGE (STGELE) CA MICS StorageMate AFTER COMPLEX SAS V8.2 Analysis of DF/SMS Space Utilization Report Date: 06/25/2004												
			Age (Days)			Allocated Space (MB)			Weighted Allocated Space (MB)			
			No. of Datasets	Minimum	Average	Maximum	Minimum	Average	Maximum	Total	Average	Total
(NONE)	DBDSALC	4	1	1,381	1,381	1,381	42	42	42	42	58,687	58,687
	MEDDSALC	1	2	21	29	37	1	7	14	14	258	515
		4	2	37	37	37	27	27	27	54	1,006	2,012
	MISCSMS	3	14	1,157	1,157	1,157	0	0	0	1	75	1,049
		4	1	713	713	713	34	34	34	34	24,239	24,239
	SMLDSALC	1	1	28	28	28	1	1	1	1	16	16
		2	1	771	771	771	1	1	1	1	437	437
		8	1	27	27	27	363	363	363	363	9,794	9,794
	TSOSTOR	1	12	36	171	281	0	0	3	4	30	361
		2	6	338	454	619	0	0	0	1	55	332
		3	1	1,448	1,448	1,448	0	0	0	0	328	328
	(OTHER)	1	1	51	51	51	5	5	5	5	249	249
		2	3	464	465	465	0	0	0	1	105	316
		4	3	335	335	335	25	25	25	76	8,541	25,624
		6	1	2,597	2,597	2,597	0	0	0	0	147	147

(sample report truncated at first page for documentation purposes)

Figure 4-59C SMS Class Usage / Analysis of DF/SMS Space Usage

SMS Implementation / By Owner (STGELC)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan

Function: Monitors conversion of data sets and volumes to DFSMS management.

Reports the amount of data converted to DFSMS management for each department/user, both in terms of data sets and amount of space.

Features: Users can optionally use VSAM and HSM data sets as part of the analysis. Users can assign data set ownership using one or more of the account fields, one or more nodes of the data set name, or using a user-defined table that translates the data set name prefix into an owner description.

As data sets are converted to DFSMS, a storage class name is assigned to each data set or set to blank for unmanaged data sets. This report accumulates totals based on the value of class name and reports totals by data set owner in terms of numbers of data sets, and number of megabytes of space allocated.

A total line reports these statistics for all data sets included in the analysis, thus providing a global view for each system.

SMS

Users: This report is useful for monitoring the DFSMS conversion effort being made by individual departments/groups/users. Running this report on a regular basis can allow you to determine the speed of your migration efforts.

COLUMN DESCRIPTIONS

Data Owner

The value used as the summarization key for this report. Depending on the values specified on the report option screen, this will represent one or more VCA account elements, one or more nodes of the data set name, or will contain a value derived from your Data Set Group table.

Total Data Sets The total number of data sets belonging to the specified owner that were included in this analysis. Depending on the options specified on the report options screen, VSAM data sets may or may not have been included.

Managed Data Sets

The total number of data sets found in the input data for this owner that were found to be managed by SMS. Depending on the SMS implementation method chosen, applications may be either fully converted (in which case, all data sets should be managed), not converted at all, or partially converted.

Unmanaged Data Sets

The total number of data sets found in the input data for this owner that were not managed by SMS. Depending on the SMS implementation method chosen, applications may be either fully converted (in which case, all data sets should be managed), not converted at all, or partially converted.

Percent of Data Sets Managed

The percentage of data sets found in the input data for this owner that were found to be managed by SMS. This is calculated by dividing the number of managed data sets by the total number of data sets found for this user.

Total Megabytes

The total amount of space allocated by the data sets found in the input data for this owner, expressed in millions of bytes (megabytes).

Managed Megabytes

The amount of space allocated by this owner belonging to data sets that were found to be SMS managed, expressed in millions of bytes (megabytes).

Unmanaged Megabytes

The amount of space allocated by this owner belonging to data sets that are not managed by SMS, expressed in millions of bytes (megabytes).

Percent of Megabytes Managed

The percentage of space allocated by this owner

belonging to data sets that are managed by SMS.
This is calculated by dividing the amount of
managed space by the total amount of space
allocated by the owner.

SMS IMPLEMENTATION (STGELC)										3
CA MICS StorageMate - OWNER REPORT										
ABC CORPORATION										
----- System Identifier=SYS7 -----										
Data Owner	Total Data Sets	Managed Data Sets	Unmanaged Data Sets	Percent of Data Sets Managed	Total Mega-Bytes	Managed Mega-Bytes	Unmanaged Mega-Bytes	Percent of Megabytes Managed		
CUSTOMER-SUPPORT	641	76	565	12	2171.7	329.1	1842.7	15		
DEVELOPMENT	8304	1856	6448	22	12858.8	2303.9	10554.9	18		
DISTRIBUTION-LIBS	107	6	101	6	662.0	1.9	660.1	0		
MARKETING-SALES	5516	1395	4121	25	19310.2	3672.0	15638.2	19		
SYSTEM-DATA	1513	34	1479	2	6279.7	20.3	6259.4	0		
-----	-----	-----	-----		-----	-----	-----			
SYSID	16081	3367	12714		41282.4	6327.2	34955.2			
	=====	=====	=====		=====	=====	=====			
	16081	3367	12714		41282.4	6327.2	34955.2			

Figure 4-60. SMS Implementation / by Owner

SMS Implementation / by Volume (STGELC)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
HSM MIG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan

Function: Monitors the implementation of DFSMS at the volume and system level.

Indicates which volumes are converted, are in the process of being converted, or are still not managed by DFSMS.

Features: Users can optionally indicate that VSAM data sets are to be included as part of the analysis. Data residing on HSM storage may also be included, and will be represented as residing on a logical volume for migration level 1, migration level 2, and HSM backup storage.

As data sets are converted to DFSMS, a non-blank value is assigned as the storage class for managed data sets. This report accumulates totals based on the value of this class and reports totals in terms of the number of data sets and megabytes allocated.

A total line reports these statistics for all volumes included in the analysis, thus providing a global view for each system.

SMS

Issues: This report is useful for monitoring the DFSMS conversion process on a volume and system level. Running this report on a regular basis can allow you to determine the speed of your migration efforts.

COLUMN DESCRIPTIONS

Volume Serial Number

The volume serial number of the storage medium. HSM data sets included in the analysis will be shown with the logical values of HSMLV1 (migration level 1), HSMLV2 (migration level 2) and HSBBAC (HSM backup).

Total Data Sets

The total number of data sets residing on the specified volume that were included in this analysis.

Managed Data Sets

The total number of data sets found in the input data for this volume that were found to be managed by SMS. Depending on the SMS implementation method chosen, volumes may be either fully converted (in which case, all data sets should be managed), not converted at all, or partially converted (known as INITIAL status). For partially converted volumes, some data sets may be managed, while some are not.

Unmanaged Data Sets

The total number of data sets found in the input data for this volume that were not managed by SMS. Depending on the SMS implementation method chosen, volumes may be either fully converted (in which case, all data sets should be managed), not converted at all, or partially converted (known as INITIAL status). For partially converted volumes, some data sets may be managed, while some are not.

Percent of Data Sets Managed

The percentage of data sets found in the input data for this volume that were found to be managed by SMS. This is calculated by dividing the number of managed data sets by the total number of data sets found for the volume.

Total Megabytes

The total amount of space allocated by the data sets found in the input data for this volume, expressed in millions of bytes (megabytes).

Managed Megabytes

The amount of allocated space on this volume belonging to data sets that were found to be SMS managed, expressed in millions of bytes (megabytes).

Unmanaged Megabytes

The amount of allocated space on this volume belonging to data sets that are not managed by SMS, expressed in millions of bytes (megabytes).

Percent of Megabytes Managed

The percentage of allocated space on this volume belonging to data sets that are managed by SMS.

This is calculated by dividing the amount of managed space by the total amount of allocated space on the volume.

Plot of Megabytes Managed

This plot represents the same value as found in the previous column (Percent of Megabytes Managed), but expressed as a plot. Each asterisk on the plot represents approximately 4% of the allocated space on the volume that is managed by SMS.

SMS IMPLEMENTATION (STGELC)										1
CA MICS StorageMate - VOLUME REPORT										
ABC CORPORATION										
----- System Identifier=SYS7 -----										
Volume Serial Number	Total Data Sets	Managed Data Sets	Unmanaged Data Sets	Percent of Data Sets Managed	Total Mega- Bytes	Managed Mega- Bytes	Unmanaged Mega- Bytes	Percent of Megabytes Managed	Plot of Megabytes Managed (Each * = 4%)	
AAA800	564	307	257	54	619.2	282.3	336.9	46	I****2****4*...6...8....I	
AAA801	117	60	57	51	612.2	279.6	332.6	46	I****2****4*...6...8....I	
AAA803	337	178	159	53	549.2	335.0	214.2	61	I****2****4*****...8....I	
AAA804	352	190	162	54	593.0	336.0	257.1	57	I****2****4*****6...8....I	
AAA806	338	155	183	46	597.4	248.9	348.5	42	I****2*****...6...8....I	
AAA807	310	165	145	53	620.0	373.5	246.5	60	I****2****4*****...8....I	
AAA900	504	274	230	54	607.0	321.3	285.7	53	I****2****4***,6...8....I	
AAA901	555	316	239	57	581.4	291.6	289.7	50	I****2****4***,6...8....I	
AAA902	348	211	137	61	603.4	308.8	294.6	51	I****2****4***,6...8....I	
AAA903	513	299	214	58	620.3	371.8	248.5	60	I****2****4*****...8....I	
AAA904	245	153	92	62	586.7	306.8	279.9	52	I****2****4***,6...8....I	
AAA905	332	204	128	61	567.5	355.2	212.2	63	I****2****4*****6*...8....I	
AAA906	201	116	85	58	609.9	314.4	295.5	52	I****2****4***,6...8....I	
AAA907	197	126	71	64	604.6	302.9	301.6	50	I****2****4***,6...8....I	
AAA908	339	191	148	56	585.2	259.5	325.7	44	I****2****4*...6...8....I	
AAA909	149	83	66	56	571.3	379.6	191.7	66	I****2****4*****6*...8....I	
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	
SYSID	5401	3028	2373		9528.1	5067.2	4460.9			
	=====	=====	=====		=====	=====	=====			
	5401	3028	2373		9528.1	5067.2	4460.9			

figure 4-61. SMS Implementation / by Volume

Space Allocation by User / Chart (STGERC)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
VCAVOA file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan

Function: Visually summarizes how each DASD volume is allocated. Chart summarization is done using one or more of the account fields, one or more nodes of the data set name, or using a user-defined table that translates the data set name prefix into an owner description. The chart makes it easy to quickly identify the major users of each volume.

Users planning for storage products that support volume pooling can use this report to detect allocation patterns useful when defining volume groups.

Features: All current VCA files are combined for this chart. The VCADAA file provides non-VSAM data set information, the VCA_VS file provides VSAM data set information, and the VCAVOA file provides free space information for each volume. Allocation information from HSM storage may also be included. The ability to select the summarization element allows users to have many different views of the same data.

Detects potential allocation problems at the volume level, and identifies groups that are allocating data sets on the wrong volumes.

Both color graphics and printer charts are supported.

SMS

Issues: Identifying the primary users of volume space is an important step when planning SMS storage groups. This report will enable you to group volumes with similar allocation patterns together, thus avoiding potential performance and capacity problems caused by conflicting users.

COLUMN DESCRIPTIONS

Volume Serial

The volume serial number of the storage medium.
Data on HSM storage will be shown with the logical values of 'HSM LV1' (migration level 1), 'HSM LV2' (migration level 2) or 'HSM BAC' (HSM backup storage).

Pie Slices

The pie slices represent the total amount of space on the volume allocated to each user, expressed in millions of bytes (megabytes), and also as a percentage of the total.

Pie Labels

The labels on each slice identify the user of the volume. Based on values coded on the report options panel, this identification can be built using one or more nodes of the data set name, one or more of the VCA account fields, or using an external Data Set Name Table to associate a data set with its owner. A value of '====> FREE SPACE' is used to represent the portion of the volume that is not assigned to any user, and is available for new allocations. Free space will not be shown for HSM storage.

Space Allocation by User / Detail (STGERC)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
VCAVOA file at the DETAIL timespan
HSMIMG file at the DETAIL timespan
HSBBAC file at the DETAIL timespan

Function: Concisely summarizes how each DASD volume is allocated. This summarization is done using one or more of the account fields, one node or more of the data set name, or using a user-defined table that translates the data set name prefix into an owner description.

Users planning for products that support volume pooling can use this report to detect allocation patterns useful in defining volume groups.

Features: All current VCA files are combined for this report. The VCADAA file provides non-VSAM data set information, the VCA_VS file provides VSAM data set information, and the VCAVOA file provides free space information for each volume. Information relating to HSM storage may also be included. The ability to select the summarization element allows users to have different views of the same data.

Detects potential allocation problems at the volume level, and identifies groups that are allocating data sets on the wrong volumes.

Lists all groups who have data sets allocated on each volume, with the groups occupying the most space listed first.

SMS

Issues: Identifying the primary users of volume space is an important step when planning SMS storage groups. This report will enable you to group volumes with similar allocation patterns together, thus avoiding potential performance and capacity problems caused by conflicting users.

COLUMN DESCRIPTIONS

Volume Serial

The volume serial number of the storage medium.

HSM storage data will be shown with the logical values of 'HSM LV1' (migration level 1), 'HSM LV2' (migration level 2) or 'HSM BAC' (HSM backup storage).

User

The identification of this user of the volume. Based on values coded on the report options panel, this identification can be built using one or more nodes of the data set name, one or more of the VCA account fields, or using an external Data Set Name Table to associate a data set with its owner. A value of '====> FREE SPACE <====' under this column is used to represent unallocated space on the volume. Free space will not be shown for HSM data.

Total # of Data Sets Present

The number of data sets allocated on this volume that belong to this user. This will contain a value of zeros for the Free Space entry.

Percent of all Data Sets

The percentage of all the data sets allocated on this volume that belong to this user. This will contain a value of zeros for the Free Space entry.

Space Allocated in MB

The total amount of space, expressed in millions of bytes (megabytes) allocated from this volume on behalf of this user. For the Free Space entry, this represents the amount of space on the volume that is unused. Entries under this column are sorted so that the primary users of the volume appear first.

Percent of all Space

The percentage of this volume that is allocated to data sets belonging to this user. For the Free Space entry, this represents the percentage of the volume that is unused. Entries under this column are sorted so that the primary users of the volume appear first.

Minimum Size (KB) Data Set

The smallest data set allocated on this volume that belongs to this user, expressed in thousands of bytes (kilobytes). This will contain a value of zeros for the Free Space entry. This will also

contain zeros for users that have data sets with VTOC entries but no tracks allocated.

Maximum Size (KB) Data Set

The largest data set allocated on this volume that belongs to this user, expressed in thousands of bytes (kilobytes). This will contain a value of zeros for the Free Space entry.

SPACE ALLOCATION BY USER (STGERC)							1
CA MICS StorageMate							
ABC CORPORATION							
----- System Identifier=SYS7 Volume Serial=LNR907 -----							
User	Total # of Data Sets Present	Percent of all Data Sets	Space Allocated in MB	Percent of all Space	Minimum Size (KB) Data Set	Maximum Size (KB) Data Set	
DEVELOPMENT	113	57.1	362.6	57.5	0	51938	
MARKETING-SALES	73	36.9	232.4	36.9	0	35605	
====> FREE SPACE <====	0	0.0	22.0	3.5	0	0	
SYSTEM-DATA	5	2.5	5.1	0.8	190	2753	
CUSTOMER-SUPPORT	4	2.0	4.2	0.7	665	2136	
DISTRIBUTION-LIBS	3	1.5	1.0	0.2	142	712	
-----	-----	-----	-----	-----	-----	-----	
VOLSER	198						
----- System Identifier=SYS7 Volume Serial=LNR908 -----							
User	Total # of Data Sets Present	Percent of all Data Sets	Space Allocated in MB	Percent of all Space	Minimum Size (KB) Data Set	Maximum Size (KB) Data Set	
MARKETING-SALES	117	33.9	312.0	49.5	0	64090	
DEVELOPMENT	179	51.9	232.1	36.8	47	28484	
====> FREE SPACE <====	0	0.0	36.0	5.7	0	0	
CUSTOMER-SUPPORT	39	11.3	20.2	3.2	47	2848	
DISTRIBUTION-LIBS	5	1.4	19.7	3.1	47	15667	
SYSTEM-DATA	7	2.0	7.3	1.2	47	4273	
-----	-----	-----	-----	-----	-----	-----	
VOLSER	347						
SYSID	545						
	=====						
	545						

Figure 4-63. Space Allocation by User / Detail

Space Management Function Simulation / Chart (STGEJB)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan

Function: Graphically represents the effect of certain space management actions on DASD volumes.

Many automated storage management software products, such as DFHSM, allow users to delete or migrate older data sets, and to release unused space from data sets that have been allocated incorrectly. This report allows users to simulate these actions and then tune these threshold values so that the desired result will be achieved.

Any combination of these different actions can be simulated--automatic deletion, automatic migration, and automatic release of over-allocated space.

Features: Users installing automated storage management products can choose thresholds that allow them to start using their new software sooner, without having an adverse effect on data availability.

Users already using automated products can use the charts to set new thresholds, or to ensure that their existing thresholds are effective.

As a user option, VSAM files can be included in the analysis.

This chart would probably be run after appropriate thresholds were defined, to provide a global view of the effect of these values on each volume.

SMS

Issues: This chart can be used to simulate most of the actions that are provided through the SMS management class. This simulation will ensure the values you specify in your management class definition will achieve the desired result. After implementation, the report can be used to monitor the effect of proposed changes to the management class definition before actually making the change.

COLUMN DESCRIPTIONS

Volume Serial Number

The volume serial number of the storage medium.

Pie Slices

The entire pie represents the number of allocated tracks on the volume that were candidates for simulation. This excludes tracks allocated by data sets that were excluded through data filtering options, as well as tracks that were unallocated on the volume. Each piece of the pie represents that portion of tracks (both in number and percentage of total) that would have been affected by a particular simulation action.

Pie Labels

The section of the pie with the label 'ALLOCATED', represents those tracks that were not affected by the simulation, and still remain allocated on the volume after the simulation completed. The other slices represent the portion of tracks on the volume that would have been freed due to migration, deletion, or partial space release.




Figure 4-64. Space Management Function Simulation / Chart

Space Management Function Simulation / Summary (STGEJB)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan

Function: Summarizes the effect of certain space management actions upon your DASD volumes.

Many automated storage management software products, such as DFHSM, allow users to delete or migrate older data sets, and to release unused space from data sets that have been allocated incorrectly. This report allows users to simulate these actions and then tune these threshold values so that the desired result will be achieved.

Any combination of these different actions can be simulated--automatic deletion, automatic migration, and automatic release of over-allocated space.

Features: Users installing automated storage management products can choose thresholds that allow them to start using their new software sooner, without having an adverse effect on data availability.

Users already using automated products can use the charts to set new thresholds, or to ensure that their existing thresholds are effective.

As a user option, VSAM files can be included in the analysis.

This report would probably be used to approximate and test the proposed threshold values, and then more detailed reports would be generated to double check the effects of these values.

Generation of the summary report is automatic.

SMS

Issues: This report can be used to simulate most of the actions that are provided through the SMS management class. This simulation will ensure the values you specify in your management class definition will achieve the desired result. After implementation, the report can be used to monitor the effect of proposed changes to the management class definition, before actually making the

change.

COLUMN DESCRIPTIONS

Volume Serial Number

The volume serial number of the storage medium.

Before # DSNs

The number of data sets on this volume that were considered candidates for simulation. Any data sets excluded by filtering criteria on the report option screen will not be included in this number. These will be referred to through the rest of this description as the candidate data sets.

Before # Trks

The number of tracks on the volume that were allocated by the candidate data sets. Any data sets excluded by filtering criteria on the report option screen will not be represented by this number. These will be referred to through the rest of this description as the candidate tracks.

Deletion Actions - # Data Sets

The number of data sets on this volume that would have been deleted using the deletion thresholds you supplied on the report option screen.

Deletion Actions - % Data Sets

The percentage of candidate data sets on this volume that would have been deleted using the deletion thresholds you supplied on the report option screen. This is calculated by dividing the number of potentially deleted data sets by the number of candidate data sets.

Deletion Actions - # Tracks

The number of tracks on this volume that would have been freed due to the simulated data set deletion.

Deletion Actions - % Tracks

The percentage of tracks on this volume that would have been freed due to the simulated data set deletion. This is calculated by dividing the number of potentially freed tracks by the number of candidate tracks.

Migration Actions - # Data Sets

The number of data sets on this volume that would have been migrated using the migration thresholds you supplied on the report option screen.

Migration Actions - % Data Sets

The percentage of candidate data sets on this volume that would have been migrated using the migration thresholds you supplied on the report option screen. This is calculated by dividing the number of potentially migrated data sets by the number of candidate data sets.

Migration Actions - # Tracks

The number of tracks on this volume that would have been freed due to the simulated data set migration.

Migration Actions - % Tracks

The percentage of tracks on this volume that would have been freed due to the simulated data set migration. This is calculated by dividing the number of potentially freed tracks by the number of candidate tracks.

Release Actions - # Data Sets

The number of data sets on this volume that would have had their excess space released using the release thresholds you supplied on the report option screen.

Release Actions - % Data Sets

The percentage of candidate data sets on this volume that would have had space released using the release thresholds you supplied on the report option screen. This is calculated by dividing the number of potentially released data sets by the number of candidate data sets.

Release Actions - # Tracks

The number of tracks on this volume that would have been freed due to the simulated partial space release.

Release Actions - % Tracks

The percentage of tracks on this volume that would have been freed due to the simulated partial space release. This is calculated by dividing the number of potentially freed tracks by the number

of candidate tracks.

Total Space Freed - # Data Sets

The number of data sets on this volume that would have been altered due to the simulation thresholds you supplied on the report option screen. This column should reflect the total number of data sets that would have been deleted, migrated, or had a portion of their space released.

Total Space Freed - % Data Sets

The percentage of candidate data sets on this volume that would have been modified using the simulation thresholds you supplied on the report option screen. This is calculated by dividing the number of potentially altered data sets by the number of candidate data sets.

Total Space Freed - # Tracks

The number of tracks on this volume that would have been freed due to the simulated actions you specified, either through data set deletion, migration or partial space release.

Total Space Freed - % Tracks

The percentage of tracks on this volume that would have been freed due to the simulated deletion, migration and release actions. This is calculated by dividing the number of potentially freed tracks by the number of candidate tracks.

SPACE MANAGEMENT FUNCTION SIMULATION (STGEJB)																		1				
CA MICS StorageMate - SUMMARY																						
ABC CORPORATION																						
----- System Identifier=SYS7 -----																						
Volume	+-	BEFORE	+-	DELETION ACTIONS				+-	MIGRATION ACTIONS				+-	RELEASE ACTIONS				+-	TOTAL SPACE FREED			
Serial	#		#	Data Sets		Tracks			Data Sets		Tracks			Data Sets		Tracks			Data Sets		Tracks	
Number	DSNs	Trks	#	%	#	%		#	%	#	%		#	%	#	%	#	%	#	%	#	%
LNR900	504	12785	17	3	390	3		99	20	1855	15		13	3	57	0	129	26	2302	18		
LNR901	555	12246	88	9	1087	9		181	33	1625	13		19	3	59	0	288	52	2771	23		
LNR902	348	12710	13	4	504	4		90	26	3017	24		13	4	54	0	116	33	3575	28		

SYSID	1407	37741	118	8	1981	5		370	26	6497	17		45	3	170	0	533	38	8648	23		

Figure 4-65. Space Management Function Simulation / Summary

Space Management Function Simulation / Detail (STGEJB)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan

Function: Provides a detailed account of the effect of certain space management actions on individual data sets and volumes.

Many automated storage management software products, such as DFHSM, allow users to delete or migrate older data sets, and to release unused space from data sets that have been allocated incorrectly. This report allows users to simulate these actions and then tune these threshold values so that the desired result will be achieved.

Any combination of these different actions can be simulated--automatic deletion, automatic migration, and automatic release of over-allocated space.

Features: Users installing automated storage management products can choose thresholds that allow them to start using their new software sooner, without having an adverse effect on data availability.

Users already using automated products can use the charts to set new thresholds, or to ensure that their existing thresholds are effective.

As a user option, VSAM files can be included in the analysis.

This detailed report would probably be run after approximate thresholds were defined, to double check the effect of those values on a data set by data set basis.

SMS

Issues: This report can be used to simulate most of the actions that are provided through the SMS management class. This simulation will ensure the values you specify in your management class definition will achieve the desired result. After implementation, the report can be used to monitor the effect of proposed changes to the management class definition, before actually making the change.

COLUMN DESCRIPTIONS

Data Set Name

The name of the data set.

Action Taken

The simulated action that would have been taken against this data set. This indicates if the data set would have been MIGRATED, DELETED, or had excess space RELEASED.

Date Created

The date that the data set was first allocated to a DASD volume.

Date Expires

The date a data set can be deleted or overwritten by another data set.

Date Last Used

The date that a data set was last referenced.

Days Since Create

The number of days between the date the data set was created and the date the data was collected.

Days Since Used

The number of days between the date the data set was last used and the date the data was collected.

Trks Freed

The number of tracks formerly occupied by this data set that would have been freed by the simulation actions. For migrated or deletion data sets, this will be the same value as tracks allocated, but for released data sets it will indicate the number of tracks that would have been released.

Trks Alloc

The number of tracks allocated to this data set on the current volume.

Trks Used

The number of tracks actually in use by this data set on the current volume.

Space Alloc (KB)

The amount of space allocated by the data set,
expressed in thousands of bytes (kilobytes).

Data Set Org

The organization of data set.

SPACE MANAGEMENT FUNCTION SIMULATION (STGEJB)											17
CA MICS StorageMate - DETAIL											
ABC CORPORATION											
----- System Identifier=SYS7 Volume Serial Number=DXB902 -----											
Data Set Name	Action Taken	Date Created	Date Expires	Date Last Used	Days Since Create	Days Since Used	Trks Freed	Trks Alloc	Trks Used	Space Alloc (KB)	Data Set Org
FFF007.PDFV2.ISPPROF	MIGRATED	16SEP87	.	14MAY90	1071	100	1	1	1	47	P0
FFF046.PDFV2.ISPPROF	MIGRATED	04JAN90	.	14MAY90	230	100	1	1	1	47	P0
QCAM.NAS9004U.IMSSUS2	MIGRATED	23MAY90	.	23JUN90	91	60	130	130	130	6172	PS
SYS6.TRACE	MIGRATED	01MAY90	.	01MAY90	113	113	2	2	0	95	PS
SYS7.PDFV2.ISPSLIB	MIGRATED	26FEB90	.	11JUN90	177	72	45	45	41	2136	P0
PROMON.R521.CNTL	MIGRATED	30MAY90	.	30MAY90	84	84	14	14	14	665	P0
PRODUCT.R521.PRODUCT.SMPMCS0	MIGRATED	06JUN90	.	14JUN90	77	69	2	2	2	95	PS
CARTGT.BSGG.ADJUST.Y002C003	RELEASED	12JUL90	.	24JUL90	41	29	2	60	1	2848	PS
CARTGT.BSGG.DISK.DET.DATA.Y002C003	RELEASED	12JUL90	.	25JUL90	41	28	3	150	75	7121	DA
LNR000.TCPCS.MASTCMEM	RELEASED	20MAR89	.	21AUG90	520	1	1	326	306	15477	DA
LNR036.DAYTIME.SOURCE	RELEASED	07AUG90	.	21AUG90	15	1	3	82	8	3893	P0
LNR063.PRIVATE.DETAIL	RELEASED	12MAR90	.	17AUG90	163	5	8	630	469	29906	DA
LNR153.FEDX.BEFORE.WEEKS	RELEASED	21AUG90	.	21AUG90	1	1	19	400	2	18988	DA
LNR182.CSS.REPORT	RELEASED	16JUL90	.	21AUG90	37	1	1	20	0	950	PS
MONDAY.S6P.WEEKS	RELEASED	09MAY90	.	08AUG90	105	14	3	200	129	9495	DA
MARKETS.PR09004U.WEEKS	RELEASED	23APR90	.	18JUL90	121	35	4	159	65	7549	DA
QCAM.INT9004C.LOAD	RELEASED	22MAY90	.	14AUG90	92	8	2	53	11	2516	P0
QCAM.SEE4160Z.AFTER1.DETAIL	RELEASED	07AUG90	.	17AUG90	15	5	2	100	47	4748	DA
QCAM.TMP9004C.SOURCE	RELEASED	04APR90	.	22AUG90	140	0	5	765	655	36316	P0
QCAM.DIC.TEXT	RELEASED	18JUN90	.	15AUG90	65	7	1	41	6	1946	P0
-----							3575	6507	4515	308913	
VOLSER							8648	16601	10764	788116	
SYSID							=====	=====	=====	=====	
							8648	16601	10764	788116	
N=	116										

Figure 4-66. Space Management Function Simulation / Detail

Tape-to-DASD Candidates (STGERB)

Source: Non-CA MICS tape management catalog

Function: Identifies data sets that are not stored on the most efficient storage media. For example, small tape data sets that are used frequently might be more cost effective if stored on DASD.

Additionally, this report allows the user to specify size and usage criteria, plus the amount of data to be moved, and then indicates the tape data sets that may be candidates for movement to DASD.

Features: Users can limit the number of data sets appearing on the report, either in terms of the megabytes represented or the number of data sets.

Users can also create a file containing the names of all candidate data sets selected. This file can be modified by the user to create JCL and control statements to move the data sets.

Reads data from an external SAS file that describes the tape management catalog. A sample external interface helps the user put their tape data into the required format.

SMS

Users: As part of SMS implementation planning, many users are identifying data sets that should be moved to a more efficient location, based on the size and usage pattern of the data set. This report will help perform that analysis. The name file that can be produced could be used to build JCL to move the data set, or could be incorporated into an ACS exit routine to allow SMS to automatically create new data sets having a matching name or the correct medium.

COLUMN DESCRIPTIONS

Data Set Name

The name of the data set. Data sets appear on this report based on their recommendation values, with the best candidates for media movement appearing first. The recommendation value is calculated as the size of the data set, multiplied

by the time since the data set was last used. The best candidates for Tape-to-DASD movement are those data sets that are small in size and referenced frequently. The number of data sets that will appear on the report is controlled by values entered on the report options screen.

Data Set Size In Megabytes

The amount of space allocated by this data set, expressed in millions of bytes (megabytes).

Volume Serial Number

The volume serial number of the storage medium.

Days Since Used

The number of days since the data set has been used. This is calculated by subtracting the date of last use from the date the data was collected.

Date Last Used

The date that the data set was last referenced.

Jobname Last Using

The name of the batch job that last used this data set.

Date Created

The date that the data set was first allocated to a DASD volume.

Date Expires

The date when a data set can be deleted or overwritten by another data set.

Block Size

The maximum length of a physical block written into this data set.

Record Length

The maximum length of each logical record in this data set.

TAPE-TO-DASD CANDIDATES (STGERB)										1
CA MICS StorageMate										
ABC CORPORATION										
Data Set Name	Data Set Size In Megabytes	Volume Serial Number	Days Since Used	Date Last Used	Jobname Last Using	Date Created	Date Expires	Block Size	Record Length	
JTK.LDM3955.FILE1	0.2	I04693	11	10AUG90	DXB999BY	10AUG90	18NOV90	80	80	
JTK.MLS2045.FILE1	0.2	I04704	13	08AUG90	DXB999BY	08AUG90	16NOV90	80	80	
JTK.AAR6565.FILE1	0.2	I04705	13	08AUG90	DXB999BY	08AUG90	16NOV90	80	80	
LNRO52.MAMICSTA.BACKUP.CHECKPT.G00489V00	20.2	M00466	1	20AUG90	DXB052DU	20AUG90	21JAN31	200	200	
LNRO52.MAMICSTI.BACKUP.CHECKPT.G0013V00	20.2	M01272	1	20AUG90	DXB052BI	20AUG90	21JAN31	200	200	
LNRO52.MAMICSTI.BACKUP.CHECKPT.G0014V00	20.2	K04142	0	21AUG90	DXB052BI	21AUG90	21JAN31	200	200	
MMAHAT.TP.BACKUP.CHECKPT.G0059V00	20.2	K01492	0	21AUG90	MSSERSBA	21AUG90	07SEP90	200	200	
QCAM.INT9004W.BACKUP.CHECKPT.G0010V00	20.2	K00853	2	19AUG90	DXB035DF	19AUG90	21AUG90	200	200	
MDG.PRODUCT.LIBLOAD.CNTL	25.6	I04735	1	20AUG90	DXB999TC	20AUG90	21FEB92	800	80	
MDG.PRODUCT.LIBLOAD.CNTL	25.6	I04736	0	21AUG90	DXB999TC	21AUG90	22FEB92	800	80	
MDG.PRODUCT.LIBLOAD.CNTL	33.6	I04733	1	20AUG90	DXB999TC	20AUG90	21FEB92	800	80	
MDG.PRODUCT.LIBLOAD.CNTL	33.6	I04734	1	20AUG90	DXB999TC	20AUG90	21FEB92	800	80	
MDG.PRODUCT.LIBLOAD.CNTL	33.6	I04742	0	21AUG90	DXB999TC	21AUG90	22FEB92	800	80	
LNRO52.MAMICSTA.BACKUP.CHECKPT.G00483V00	20.2	K00225	10	11AUG90	DXB052DU	11AUG90	21JAN31	200	200	
LNRO52.MAMICSTA.BACKUP.CHECKPT.G00484V00	20.2	K04096	9	12AUG90	DXB052DU	12AUG90	21JAN31	200	200	
LNRO52.MAMICSTA.BACKUP.CHECKPT.G00485V00	20.2	K03676	8	13AUG90	DXB052DU	13AUG90	21JAN31	200	200	
LNRO52.MAMICSTA.BACKUP.CHECKPT.G00486V00	20.2	K03515	7	14AUG90	DXB052DU	14AUG90	21JAN31	200	200	
LNRO52.MAMICSTA.BACKUP.CHECKPT.G00487V00	20.2	K01485	6	15AUG90	DXB052DU	15AUG90	21JAN31	200	200	
LNRO52.MAMICSTA.BACKUP.CHECKPT.G00488V00	20.2	K01648	4	17AUG90	DXB052DU	17AUG90	21JAN31	200	200	
LNRO52.MAMICSTI.BACKUP.CHECKPT.G0007V00	20.2	K04100	9	12AUG90	DXB052BI	12AUG90	21JAN31	200	200	
LNRO52.MAMICSTI.BACKUP.CHECKPT.G0008V00	20.2	K00226	8	13AUG90	DXB052BI	13AUG90	21JAN31	200	200	
LNRO52.MAMICSTI.BACKUP.CHECKPT.G0009V00	20.2	K01097	7	14AUG90	DXB052BI	14AUG90	21JAN31	200	200	
LNRO52.MAMICSTI.BACKUP.CHECKPT.G0010V00	20.2	K00200	6	15AUG90	DXB052BI	15AUG90	21JAN31	200	200	
LNRO52.MAMICSTI.BACKUP.CHECKPT.G0011V00	20.2	M00197	5	16AUG90	DXB052BI	16AUG90	21JAN31	200	200	
LNRO52.MAMICSTI.BACKUP.CHECKPT.G0012V00	20.2	K01655	4	17AUG90	DXB052BI	17AUG90	21JAN31	200	200	
MMAHAT.TP.BACKUP.CHECKPT.G0056V00	20.2	K03670	5	16AUG90	MSSERSBA	16AUG90	02SEP90	200	200	
OSEHJH.SRLT.BACKUP.CHECKPT.G0004V00	20.2	K00448	8	13AUG90	DXB145BA	13AUG90	21JAN31	200	200	
QCAM.INT9004W.BACKUP.CHECKPT.G0001V00	20.2	K04251	5	16AUG90	DXB035RF	05AUG90	21AUG90	200	200	
QCAM.INT9004W.BACKUP.CHECKPT.G0006V00	20.2	K03277	5	16AUG90	DXB035RF	09AUG90	21AUG90	200	200	
QCAM.INT9004W.BACKUP.CHECKPT.G0007V00	20.2	K04232	5	16AUG90	DXB035RF	14AUG90	21AUG90	200	200	
QCAM.INT9004W.BACKUP.CHECKPT.G0008V00	20.2	K03065	5	16AUG90	DXB035DF	16AUG90	21AUG90	200	200	
QCAM.INT9004W.BACKUP.CHECKPT.G0009V00	20.2	K01184	5	16AUG90	DXB035DF	16AUG90	21AUG90	200	200	
QCAM.INT9004W.MBACKUP.CHECKPT.G0002V00	20.2	K04266	5	16AUG90	DXB035RE	15AUG90	21AUG90	200	200	
QCAM.INT9004W.MBACKUP.CHECKPT.G0003V00	20.2	K01167	4	17AUG90	DXB035RF	16AUG90	21AUG90	200	200	
QCAM.TST8910T.BACKUP.CHECKPT.G0005V00	20.2	M02874	5	16AUG90	DXB035RE	25JUL90	21JAN31	200	200	
QCAM.XYZ9004R.BACKUP.CHECKPT.G0002V00	20.2	K04220	8	13AUG90	DXB225RE	08AUG90	21JAN31	200	200	
QCAM.XYZ9004R.BACKUP.CHECKPT.G0004V00	20.2	M00624	8	13AUG90	DXB225RE	10AUG90	21JAN31	200	200	
JTK.PRODUCT.LIBLOAD.CNTL	25.6	I04730	4	17AUG90	DXB999TC	17AUG90	18FEB92	800	80	
PRODUCT.CPU.DATA	55.4	T51800	1	20AUG90	DXB117TT	20AUG90	28NOV90	6160	80	
PRODUCT.CPU.DATA	55.4	M00433	2	19AUG90	DXB117TT	19AUG90	27NOV90	6160	80	
=====										
874.8										
N=	40									

Figure 4-67. Tape-to-DASD Candidates

Unmanageable Data Sets / Summary (STGELA)

Source: VCADAA file at the DETAIL timespan

Function: Summarizes, by volume, errors that must be corrected before the volume can be used with DFSMS or other automated storage management products.

Users planning for automated products can identify those volumes with potential problems and those volumes that can be automated immediately. This shortens the time needed to implement automated products, and reduces the errors that occur when the software is implemented.

Features: For some error types, JCL and control statements that will correct the error can be generated automatically.

The error types listed on this report include older access methods, uncataloged data sets, and data sets with restrictive allocation parameters.

Selection of only certain error types is supported, so that each user can tailor the detection and correction process to their own needs.

SMS

Users: This report can be used to identify data sets that cannot be managed by SMS. These data sets will then either have to be converted to a supported format or moved to a volume that will not be managed by SMS. The JCL and control statements generated by the analysis can assist in correcting some types of errors.

COLUMN DESCRIPTIONS

Volume Serial Number

The volume serial number of the storage medium.

TOTAL ERRORS

The total number of unmanageable data sets found on this volume. This column is a sum of all the error columns to the right of it.

Catalog Errors

The number of data sets found on this volume that

either were not cataloged, or had a catalog entry showing the data set residing on a different volume from this one. Catalog errors can only be detected if the data being analyzed was produced by a run of the CA MICS DASD Collector that had the catalog option active.

Zero Extents

The number of data sets on this volume that were represented by an entry in the volume table of contents (VTOC), but had no space allocated on the volume. Review these data sets to see if they are really necessary.

Unmoveable Data Sets

The number of data sets on this volume that had the unmoveable attribute set in the data set organization option area, or were defined as direct-access data sets with an OPTCD value of 'A'. This indicates that the data sets cannot be moved, and thus cannot be managed with automated storage management software.

Absolute Track Allocation

The number of data sets on this volume that were allocated at an absolute track location on the volume. Review these data sets to see why they were allocated in this manner. This indicates that the data sets cannot be managed with automated storage management software.

Invalid DSORG

The number of data sets on this volume that contained missing or invalid values in the data set organization option area, as well as data sets defined as using the ISAM access method. Investigate these data sets to see if they contain valid data, or if they are empty. This indicates that the data sets cannot be managed with automated storage management software.

CVOL Catalogs

The number of old-style CVOL catalogs that were found on this volume. These types of catalogs are not supported by many automated software products, nor are any of the data sets cataloged in them. Consider converting these to a more recent catalog type, such as an ICF (Integrated Catalog Facility) catalog.

VSAM Data Sets

The number of VSAM data sets on this volume that were found to be cataloged in old style-VSAM catalogs. These older types of catalogs will appear on this report, as well as any data sets found cataloged in them. These types of catalogs are not supported by many automated software products, nor are any of the data sets cataloged in them. Consider converting these to a more recent catalog type, such as an ICF (Integrated Catalog Facility) catalog.

Password Protected

The number of data sets on this volume that were found to be protected using the old-style OS password mechanism. These types of data sets are not supported by many automated software products. Consider converting these to a more recent security mechanism, such as RACF (Resource Access Control Facility) or a similar facility.

UNMANAGEABLE DATA SETS (STGELA)									
CA MICS StorageMate - SUMMARY									
ABC CORPORATION									
----- System Identifier=SYSA -----									
Volume Serial Number	* TOTAL * ERRORS *	Catalog Errors	Zero Extents	Unmoveable Data Sets	Absolute Track Allocation	Invalid DSORG	CVOL Catalogs	VSAM Data Sets	Password Protected
MVSMG1	192	188	1	0	0	3	0	0	0
MVSMG2	120	117	0	0	0	0	0	0	3
TS0053	62	49	3	0	0	10	0	0	0
TS0054	61	50	5	0	0	6	0	0	0
TS0021	52	40	7	0	0	5	0	0	0
TS0023	52	37	11	0	0	4	0	0	0
TS0029	51	46	3	0	0	2	0	0	0
TS0063	49	37	8	0	0	4	0	0	0
TS0055	48	38	2	0	0	8	0	0	0
TS0057	48	34	5	0	0	9	0	0	0
TS0027	47	37	8	0	0	1	0	0	1
TS0056	47	37	4	0	0	6	0	0	0
TS0035	46	29	11	0	0	6	0	0	0
TS0022	45	39	3	0	0	3	0	0	0
TS0031	45	31	10	0	0	4	0	0	0
TS0026	43	34	0	0	0	9	0	0	0
TS0028	43	32	5	0	0	6	0	0	0
TS0020	42	30	5	0	0	7	0	0	0
TS0034	42	29	6	0	0	6	0	0	1
TS0024	41	32	4	0	0	4	0	0	1
TS0033	41	25	10	0	0	5	0	0	1
TS0059	40	29	9	0	0	2	0	0	0
TS0065	39	34	0	0	0	5	0	0	0
TS0036	36	26	4	0	0	6	0	0	0
TS0064	34	28	0	0	0	6	0	0	0
TS0037	32	17	7	0	0	8	0	0	0
TS0030	31	25	2	0	0	4	0	0	0
TS0062	23	23	0	0	0	0	0	0	0
TS0061	16	15	0	0	0	1	0	0	0
TS0060	6	6	0	0	0	0	0	0	0
TS0025	0	0	0	0	0	0	0	0	0
TS0051	0	0	0	0	0	0	0	0	0
TS0052	0	0	0	0	0	0	0	0	0
TS0067	0	0	0	0	0	0	0	0	0
TS0068	0	0	0	0	0	0	0	0	0
SYSID	1474	1194	133	0	0	140	0	0	7

Figure 4-68. Unmanageable Data Sets / Summary

Unmanageable Data Sets / Detail (STGELA)

Source: VCADAA file at the DETAIL timespan

Function: Identifies data sets that may not be supported by automated storage management products until they are modified or converted.

Users planning for automated products can identify those volumes with potential problems and those volumes that can be automated immediately. This shortens the time needed to implement automated products, and reduces the errors that occur when the software is implemented.

Features: For some error types, JCL and control statements that will correct the error can be generated automatically.

The error types listed on this report include older access methods, uncataloged data sets, and data sets with restrictive allocation parameters.

Selection of only certain error types is supported, so that each user can tailor the detection and correction process to their own needs.

SMS

Users: This report can be used to identify data sets that cannot be managed by SMS. These data sets will then either have to be converted to a supported format or moved to a volume that will not be managed by SMS. The JCL and control statements generated by the analysis can assist in correcting some types of errors.

COLUMN DESCRIPTIONS

Data Set Name

The name of the data set.

Volume Serial Number

The volume serial number of the storage medium.

Data Set Org

The organization of the data set.

Date Created

The date that the data set was first allocated to a DASD volume.

Date Last Used

The date that the data set was last referenced.

Alloc Type

The type of space allocation for this data set.

Extents

The number of non-contiguous segments of space allocated to the data set on one volume.

Prot. Type

The security level of the data set.

Alloc Size (MB)

The amount of space allocated by the data set, expressed in millions of bytes (megabytes).

Error Condition

A description of the error that causes this data set to be unmanageable. A description of the values that will appear in this column are as follows:

UNKNOWN DSORG / POSSIBLY EMPTY

The data set organization option for this data set has never been set to a valid value. While some exotic software can still perform I/O to the data set without setting the organization option, there is also a possibility that the data set was allocated but never used.

UNMOVEABLE DATA SET

The data set organization option for this data set contains a 'U', indicating that the data set cannot be moved to another location. This prevents this data set from being managed by HSM, SMS, or other automated data management products.

UNMOVEABLE DATA SET (OPTCD=A)

The data set organization option for this data set indicates a direct-access (DA) data set, and an OPTCD value of 'A' was also specified. This means that the data contains location-sensitive information, and the

application would probably not work if the data set were moved to another location. This prevents this data set from being managed by HSM, SMS, or other automated data management products.

ABSOLUTE TRACK ALLOCATION

The allocation type option for this data set indicates that the data set was allocated at an absolute track location on the volume. This could have been done for performance reasons, or because the software using the data set expects it to reside at a fixed location. This prevents this data set from being managed by HSM, SMS, or other automated data management products.

UNIQUE VSAM (NON-ICF) DATA SET

SHARED VSAM (NON-ICF) DATA SET

A VSAM data set (either UNIQUE or SHARED) was detected that was cataloged in an old-style VSAM catalog. These types of catalogs are not supported by many automated data management products, and the data sets contained in them should be converted to use the newer ICF (Integrated Catalog Facility) catalogs. These messages may also appear for the VSAM type catalogs themselves, as they are also VSAM data sets.

CVOL CATALOG

An OS CVOL catalog was detected on this volume. These types of catalogs are not supported by many automated data management products, and the data sets contained in them should be converted to use the newer ICF (Integrated Catalog Facility) catalogs.

PASSWORD PROTECTED

This data set was found to be protected using the old-style OS password mechanism. These types of data sets are not supported by many automated software products. Consider converting the data set to a more recent security mechanism, such as RACF (Resource Access Control Facility) or a similar facility.

ISAM ACCESS METHOD

This data set is defined as using the Indexed Sequential Access Method (ISAM). This is an older access method that is not supported by many automated data management products. Consider converting the data set to a more recent access method, such as VSAM.

ZERO EXTENTS / NO SPACE USED

This data set is defined in the Volume Table of Contents (VTOC), but is using no space on the volume. While there are some valid reasons for having this type of data set, it could also be related to an allocation error or a previous requirement that is no longer valid.

NOT CATALOGED

CATALOG INFORMATION INCORRECT

This data set is either not cataloged, or has catalog information that is incorrect. Many automated data management products will only support data sets that are cataloged. The data set should either be deleted, or the catalog information for it should be updated or added. Catalog errors can only be detected if the data being analyzed was produced by a run of the CA MICS DASD Collector that had the catalog option active.

UNMANAGEABLE DATA SETS (STGELA)								
CA MICS StorageMate - DETAIL								
ABC CORPORATION								
----- System Identifier=SYSA Volume Serial Number=TS0104 -----								
Data Set Name	Data Set Org	Date Created	Date Last Used	Alloc Type	Ext-ents	Prot. Type	Alloc Size (MB)	ERROR CONDITION
APLTEST.ISPF16.ISPPROF	PO	07MAR91	07MAR91	TRK	1	NONE	0.1	NOT CATALOGED
BIGTST1.VMF32.HCLLOG	***	24JUL90	.	CYL	1	NONE	0.7	UNKNOWN DSORG / POSSIBLY EMPTY
BIGTST2.ISPF.ISPPROF	PO	05MAR91	05MAR91	TRK	1	NONE	0.1	NOT CATALOGED
BIGTST2.TEMPLIB	PS	11MAR91	.	BLK	0	NONE	0.0	ZERO EXTENTS / NO SPACE USED
CXR.TST30VXA.MSGFILE	***	15APR91	.	CYL	1	NONE	1.4	UNKNOWN DSORG / POSSIBLY EMPTY
DAVES81.ISPF.ISPPROF	PO	30OCT90	22MAR91	TRK	1	NONE	0.1	NOT CATALOGED
DBC.STDCMP.ASM	PO	11FEB87	15APR91	TRK	5	NOPR	3.0	PASSWORD PROTECTED
DBC.STDCMP.MACRO	PO	27JAN87	16APR91	TRK	11	NOPR	5.3	PASSWORD PROTECTED
DBC.DEV.AUDIT.LOG2	PS	30JAN91	16APR91	CYL	0	NONE	0.0	ZERO EXTENTS / NO SPACE USED
DEVPE22.ASXJES3.LNK	PS	11MAR91	13MAR91	TRK	0	NONE	0.0	ZERO EXTENTS / NO SPACE USED
DEVWD11.DSIWD22.D91100.T105518	***	10APR91	.	TRK	1	NONE	0.6	UNKNOWN DSORG / POSSIBLY EMPTY
DEV046.MNGT.USER.BATCH32.LOG	PS	15FEB91	24FEB91	TRK	0	NONE	0.0	ZERO EXTENTS / NO SPACE USED
DEV159.CLIST	PO	31JAN91	16APR91	TRK	1	NONE	0.1	NOT CATALOGED
DEV186.AUS.NCAL5	PS	20FEB91	25FEB91	CYL	0	NONE	0.0	ZERO EXTENTS / NO SPACE USED
ESP.LOCKFKL	***	16APR91	.	CYL	1	NONE	0.8	UNKNOWN DSORG / POSSIBLY EMPTY
FDXTB.A.TST4975.INSTALL.LIB	PO	08FEB91	08FEB91	TRK	1	NONE	0.1	NOT CATALOGED
FXAPCFR.ISPFXA.ISPPROF	PO	03APR90	16APR91	TRK	1	NONE	0.1	NOT CATALOGED
FXAPE22.SORTWK01	PS	07MAR91	08MAR91	CYL	0	NONE	0.0	ZERO EXTENTS / NO SPACE USED
FXARE11.ISPFXA.ISPPROF	PO	11DEC89	14APR91	TRK	2	NONE	0.1	NOT CATALOGED
FXASR11.SMSTEST1.DATA	***	16APR91	.	TRK	1	NONE	0.6	UNKNOWN DSORG / POSSIBLY EMPTY
FXASYS.MISSBD.USER.LOAD	PO	09APR91	16APR91	BLK	1	NONE	0.1	NOT CATALOGED
FXA042.D91057.T105846.FILE	***	26FEB91	.	TRK	1	NONE	0.6	UNKNOWN DSORG / POSSIBLY EMPTY
FXA069.MISS.USER.REPORT	PS	04MAR91	04MAR91	TRK	0	NONE	0.0	ZERO EXTENTS / NO SPACE USED
FXA084.ISPPROF	PO	11JUL89	16APR91	BLK	1	NONE	0.1	NOT CATALOGED
FXA106.VB.CLIST	PO	09SEP87	16APR91	TRK	1	NONE	0.1	NOT CATALOGED

Figure 4-69. Unmanageable Data Sets / Detail

Volume Configuration / Summary (STGEPB)

Source: VCAVOA file at the DETAIL timespan
HARDVA file at the DAYS timespan

Function: Concisely summarizes performance and capacity information at the device type level. This information is reported for each unique device type and for each logical system (SYSID).

This makes it easy to identify potential performance, capacity or balance problems at the device type level.

Users planning for storage products that support volume pooling can use this report to design volume groupings that will find combinations that provide optimum availability and free space.

Features: Two different CA MICS files are combined to produce this report; one provides space information and the other provides performance information. The HARDVA file can be filtered at the hour or zone level so that the performance information reflects a specific time period.

Volumes that appear in the HARDVA file but not the VCAVOA file (because they were not scanned) can be included or excluded by user option. If included, values for these volumes are not included in the space totals, but are included in the response time and I/O rate values.

Generation of the summary report is automatic.

SMS Users: The response time and capacity values obtained from this report can help when defining SMS storage classes and storage groups.

COLUMN DESCRIPTIONS

Device Type
The type of I/O device.

Number of Volumes
The number of volumes found within the input data that were found to have this device type.

Total Capacity (MB)

The total potential capacity of all the volumes of this device type, expressed in millions of bytes (megabytes).

Space Allocated (MB)

The total amount of space allocated by all the data sets on all the volumes of this device type, expressed in millions of bytes (megabytes).

Space Used (MB)

The total amount of space used by all the data sets on all the volumes of this device type, expressed in millions of bytes (megabytes).

Space Free (MB)

The total amount of space free (unallocated) on all the volumes of this device type, expressed in millions of bytes (megabytes).

Total Data Sets

The total number of data sets residing on all the volumes with this device type.

Non-VSAM # Data Sets

The number of non-VSAM data sets residing on all the volumes with this device type.

Non-VSAM Alloc (MB)

The total amount of space, expressed in millions of bytes (megabytes), allocated by all the non-VSAM data sets residing on volumes with this device type.

VSAM # Data Sets

The number of VSAM data sets residing on all the volumes with this device type.

VSAM Alloc (MB)

The total amount of space, expressed in millions of bytes (megabytes), allocated by all the VSAM data sets residing on volumes with this device type.

Average I/O Rate

The total number of I/O requests per second being received by all the volumes with this device type. This value is the sum of the rates for the

individual volumes within the summarization.

Average Response Time

The average time a user of this device type waits, in thousandths of a second (milliseconds), for an I/O request to complete. This is recalculated at the device type level, using the same technique used to calculate this value for an individual volume.

VOLUME CONFIGURATION (STGEPS)												5
CA MICS StorageMate - SUMMARY												
ABC CORPORATION												
----- System Identifier=SYS7 -----												
Device	Number	Total	Space	Space	Space	Total	--- Non-VSAM ---	--- VSAM ---	--- VSAM ---	--- VSAM ---	--- Average ---	
Type	of	Capacity	Allocated	Used	Free	Data	# Data	Alloc	# Data	Alloc	I/O	Response
	Volumes	(MB)	(MB)	(MB)	(MB)	Sets	Sets	(MB)	Sets	(MB)	Rate	Time
3380	78	49664.77	42616.69	27858.76	6836.06	14773	13812	36105.71	961	6510.98	2.640	0.026
3390-2	10	6367.32	6228.07	4778.48	106.97	2399	2300	5606.31	99	621.76	1.537	0.026

SYSID	88	56032.09	48844.76	32637.24	6943.03	17172	16112	41712.02	1060	7132.74	2.515	0.026

Figure 4-70. Volume Configuration / Summary

Volume Configuration / Detail (STGEPB)

Source: VCAVOA file at the DETAIL timespan
HARDVA file at the DAYS timespan

Function: Concisely summarizes performance and capacity information at the volume level. This information is reported for each volume and string.

This makes it easy to identify potential performance or capacity problems, at both the volume and string level.

Users planning for storage products that support volume pooling can use this report to design various volume groupings to find combinations that provide optimum performance and availability.

Features: Two different CA MICS files are combined to produce this report; one provides space information and the other provides performance information. The HARDVA file can be filtered at the hour or zone level so that the performance information reflects a specific time period.

Volumes that appear in the HARDVA file but not the VCAVOA file (because they were not scanned) can be excluded by user option or may still appear on the report, but will have 'N/A' or zeros in those fields obtained from the VCAVOA file.

SMS Users: The response time and capacity values obtained from this report can help when defining SMS storage classes and storage groups.

COLUMN DESCRIPTIONS

Control Unit Address
A three-byte field that gives the logical address of an I/O control unit.

Device Address
The physical address of an I/O device.

C.U. Path
The Logical Control Unit (MVS/XA users) or Logical Channel (non-XA users) through which this volume

is attached to the I/O subsystem.

Volume Serial Number

The volume serial number of the storage medium.

Device Type

The type of I/O device.

Volume Usage

The current DASD Volume Use Attributes. The value '(N/A)' will appear in this column if this volume was not found in the VCAVOA file, but was found in the HARDVA file, and the unmatched volume option was specified on the report options screen.

SMS Volume Status

The current System Managed Volume Status. The value '(N/A)' will appear in this column if this volume was not found in the VCAVOA file, but was found in the HARDVA file, and the unmatched volume option was specified on the report options screen.

Volume Free

The percent of space that is available for allocation on this volume. A value of zero will appear in this column if this volume was not found in the VCAVOA file, but was found in the HARDVA file, and the unmatched volume option was specified on the report options screen.

Fragmentation Index

A measure of how fragmented the allocated space is for the current volume. A value of zero will appear in this column if this volume was not found in the VCAVOA file, but was found in the HARDVA file, and the unmatched volume option was specified on the report options screen.

Total Data Sets

The total number of data sets residing on this volume. A value of zero will appear in this column if this volume was not found in the VCAVOA file, but was found in the HARDVA file, and the unmatched volume option was specified on the report options screen.

Non-VSAM Data Sets

The percentage of Non-VSAM data sets residing on this volume. This is calculated by dividing the

number of Non-VSAM data sets by the total number of data sets on the volume. A value of zero will appear in this column if this volume was not found in the VCAVOA file, but was found in the HARDVA file, and the unmatched volume option was specified on the report options screen.

VSAM Data Sets

The percentage of VSAM data sets residing on this volume. This is calculated by dividing the number of VSAM data sets by the total number of data sets on the volume. A value of zero will appear in this column if this volume was not found in the VCAVOA file, but was found in the HARDVA file, and the unmatched volume option was specified on the report options screen.

Total I/O Rate

The average number of I/O requests per second being received by this volume.

Average Response Time

The average time a user of this volume waits, in thousandths of a second (milliseconds), for an I/O request to complete.

VOLUME CONFIGURATION (STGE PB)													1
CA MICS StorageMate - DETAIL													
ABC CORPORATION													
----- System Identifier=SYS7 / Control Unit Address=140 -----													
Device Address	C.U. Path	Volume Serial Number	Device Type	Volume Usage	SMS Volume Status	Volume Free	Frag. Index	+----- Total	Data Non- VSAM	Sets VSAM	+----- VSAM	Total I/O Rate	Average Response Time
140	0001	WORK41	3380	PUBLIC	INITIAL	38%	0.314	60	98%	2%		4.185	0.026
141	0001	DIST02	3380	PRIVATE		16%	0.052	117	90%	10%		15.808	0.017
142	0001	PAG403	3380	PRIVATE		29%	0.099	6	17%	83%		8.657	0.038
143	0001	LNR902	3380	STORAGE		3%	0.948	351	99%	1%		2.421	0.024
144	0001	SP40L1	3380	PRIVATE		13%	0.375	134	100%	0%		0.001	0.030
145	0001	SP00L1	3380	PRIVATE		0%	0.312	2	100%	0%		10.424	0.023
146	0001	LNR915	3380	STORAGE		2%	0.463	352	98%	2%		0.812	0.020
147	0001	LNR904	3380	STORAGE		4%	0.805	248	99%	1%		1.204	0.022
						13%	0.421	1270	98%	2%		43.512	0.024
----- System Identifier=SYS7 / Control Unit Address=141 -----													
Device Address	C.U. Path	Volume Serial Number	Device Type	Volume Usage	SMS Volume Status	Volume Free	Frag. Index	+----- Total	Data Non- VSAM	Sets VSAM	+----- VSAM	Total I/O Rate	Average Response Time
148	0001	SPAR05	3380	PUBLIC		2%	0.338	139	94%	6%		0.252	0.026
149	0001	SPAR03	3380	PUBLIC		7%	0.399	64	98%	2%		3.196	0.021
14A	0001	LNR919	3380	STORAGE		0%	0.833	229	95%	5%		0.412	0.020
14B	0001	LNR906	3380	STORAGE		3%	0.793	202	100%	0%		1.111	0.023
14C	0001	LNR917	3380	STORAGE		1%	0.321	207	97%	3%		1.418	0.025
14D	0001	LNR907	3380	STORAGE		3%	0.689	198	99%	1%		1.084	0.021
14E	0001	LNR918	3380	STORAGE		0%	0.927	236	99%	1%		0.392	0.020
14F	0001	LNR908	3380	STORAGE		6%	0.803	345	98%	2%		1.793	0.019
						3%	0.638	1620	98%	2%		9.658	0.022
----- System Identifier=SYS7 / Control Unit Address=240 -----													
Device Address	C.U. Path	Volume Serial Number	Device Type	Volume Usage	SMS Volume Status	Volume Free	Frag. Index	+----- Total	Data Non- VSAM	Sets VSAM	+----- VSAM	Total I/O Rate	Average Response Time
240	0002	VIPL01	3380	PRIVATE		8%	0.337	136	96%	4%		3.156	0.027
241	0002	LNR909	3380	STORAGE		5%	0.652	159	94%	6%		0.960	0.015
242	0002	PAG401	3380	PRIVATE		44%	0.131	6	17%	83%		7.144	0.053
243	0002	LNR910	3380	STORAGE		2%	0.461	224	92%	8%		1.117	0.027
244	0002	ITS002	3380	PRIVATE		1%	0.502	576	94%	6%		0.079	0.022
245	0002	LNR911	3380	STORAGE		1%	0.382	257	98%	2%		1.827	0.032
246	0002	LNR923	3380	STORAGE		7%	0.665	213	100%	0%		0.448	0.028
247	0002	LNR912	3380	STORAGE		4%	0.259	217	98%	2%		0.184	0.025
						9%	0.424	1788	95%	5%		14.916	0.039

Figure 4-71. Volume Configuration / Detail

Analysis of I/O Rate and Response Time (STGEPB)

Source: VCAVOA file at the DETAIL timespan
HARDVA file at the DAYS timespan

Function: Concisely summarizes the output of clustering, including the number and percentage of volumes, aggregate I/O rate and response times. The results are split into two sections: "NORMAL", which contains all volumes whose performance is considered normal by statistical standards; and "EXCEPTIONAL", which includes all volumes operating outside of a normal range.

Examination of the "NORMAL" section provides a high-level overview of subsystem performance and can serve as a benchmark for additional studies. Of greater interest is the "EXCEPTIONAL" section, which identifies groups of volumes requiring further analysis.

Features: Two different CA MICS files are combined to produce this report; one provides space information and the other provides performance information. The HARDVA file can be filtered at the hour or zone level so that the performance information reflects a specific time period.

Volumes that appear in the HARDVA file but not the VCAVOA file (because they were not scanned) can be excluded by user option or may still appear on the report, but will have 'N/A' or zeros in those fields obtained from the VCAVOA file.

SMS N/A
Users:

COLUMN DESCRIPTIONS

Description

A text string indicating whether the volumes are either "NORMAL" or "EXCEPTIONAL" in their performance.

Cluster

The identifier of the cluster, an integer from one to n.

Number of Volumes

The number of volumes represented by each cluster.

Percent of Volumes

The percentage of the total storage subsystem represented by the volumes within each cluster.

Aggregate I/O Rate

The total I/O rate of all volumes within each cluster.

Percent of I/O Activity

The percentage of all I/O operations by the storage subsystem represented by the volumes within each cluster.

Average Response Time

The average (mean) response time for all volumes represented within each cluster.

Average Weighted Response Time

The mean response time weighted by the I/O rate for all volumes represented within each cluster. This metric is provided as an alternative to avoid situations where the Average Response Time value is seriously distorted by a small number of outliers. The two response time measurements ideally will be the same or very close in value. A large difference in values indicates a distortion of the Average Response Time values.

VOLUME CONFIGURATION (STGEPS) CA MICS StorageMate - SUMMARY AFTER COMPLEX SAS V8.2 Cluster Analysis of I/O Rate and Response Time							
Description	Cluster	Number of Volumes	Percent of Volumes	Aggregate I/O Rate	Percent of I/O Activity	Average Response Time	Avg. Wgtd. Response Time
Normal Volumes	6	1,641	80.64	51.8	7.56	0.004	0.003
	15	218	10.71	52.4	7.64	0.014	0.006
	14	79	3.88	73.0	10.65	0.002	0.002
	12	22	1.08	8.5	1.24	0.010	0.008
	8	8	0.39	13.2	1.93	0.001	0.001
	16	2	0.10	3.2	0.46	0.002	0.002
-----		-----		-----		-----	
Normal Volumes		1,970	96.81	202.0	29.50		
Exception Volumes	1	15	0.74	356.1	52.00	0.049	0.004
	16	10	0.49	20.4	2.98	0.002	0.002
	8	7	0.34	12.8	1.86	0.001	0.001
	9	5	0.25	13.4	1.96	0.002	0.002
	12	5	0.25	2.3	0.34	0.013	0.010
	13	5	0.25	3.3	0.48	0.011	0.009
	4	4	0.20	11.9	1.74	0.002	0.002
	3	3	0.15	9.2	1.35	0.003	0.003
	11	3	0.15	14.4	2.10	0.002	0.002
	2	2	0.10	8.1	1.19	0.003	0.003
	5	2	0.10	7.0	1.03	0.001	0.001
	7	2	0.10	16.1	2.34	0.002	0.002
	10	2	0.10	7.7	1.13	0.002	0.002
-----		-----		-----		-----	
Exception Volumes		65	3.19	482.9	70.50		
		=====	=====	=====	=====		
		2,035	100.00	684.9	100.00		

4-71A Cluster Analysis of I/O Rate and Response Time

Listing of All Exception Volumes by Cluster (STGEPB)

Source: VCAVOA file at the DETAIL timespan
 HARDVA file at the DAYS timespan

Function: A detailed listing of all volumes that are
 considered "exceptional" during the clustering
 process.

 This permits a rapid review of those volumes that
 represent potential performance problems, so that
 they can be subjected to further analysis.

Features: Two different CA MICS files are combined to
 produce this report; one provides space
 information and the other provides performance
 information. The HARDVA file can be filtered at
 the hour or zone level so that the performance
 information reflects a specific time period.

 Volumes that appear in the HARDVA file but not the
 VCAVOA file (because they were not scanned) can be
 excluded by user option or may still appear on the
 report, but will have 'N/A' or zeros in those
 fields obtained from the VCAVOA file.

SMS N/A

Users:

COLUMN DESCRIPTIONS

Cluster
The identification of the reporting cluster; an integer from one to n.

Logical Control Unit (ID)
A four byte field assigned by the z/OS IOCP process that identifies the specific logical control unit (LCU) where each volume resides. The LCU is a logical representation of one or more physical control units and their associated I/O devices and is used to control (queue) all I/O operations.

Control Unit Address
A three-byte field that gives the logical address

of an I/O control unit.

Device Number

The physical address of an I/O device.

Volume (Serial Number)

The volume serial number of the storage medium.

Device I/O Rate

The average (mean) number of I/O operations, per second for the reported volume.

Average Response Time

The average (mean) response time for the reported volume. This is calculated as the sum of device connect, disconnect and pending (busy) times, divided by the total number of I/O operations (SSCH Count). The result is then added to the average queue time.

VOLUME CONFIGURATION (STGEPB)						
CA MICS StorageMate - DETAIL						
AFTER COMPLEX SAS V8.2						
Listing of All Exception Volumes, by Cluster						
Cluster	Logical Control Unit	Control Unit Addr.	Device Number	Volume	Device I/O Rate	Average Response Time
1	003F	21D1	21DD	TS026A	0.0	0.384
	003F	2151	215A	MVCA01	71.5	0.003
	004A	2EA0	2EA5	PANS03	6.6	0.001
	004A	2EA1	2EAB	PANS09	11.0	0.001
	004A	2E91	2E9A	MVS008	46.0	0.004
	0041	23C1	23CD	TSU014	0.4	0.056
	0042	2431	243D	MVR14D	24.6	0.003
	0042	2461	246E	NMD003	6.4	0.002
	0043	2550	2554	DISK01	12.6	0.011
	0043	2571	257F	TS023A	2.4	0.019
	0044	2641	264E	TS010B	1.1	0.009
	0045	27D0	27D4	OSI006	151.9	0.001
	0049	2D70	2D70	JESCK1	2.0	0.006
	006D	4300	4301	DMP014	1.2	0.218
	006D	4300	4307	DSK110	18.5	0.012

1					356.1	
2	004A	2E20	2E25	MVS006	4.0	0.003
	0044	26C0	26C1	NMD035	4.1	0.002

2					8.1	
3	004A	2E30	2E35	RTM005	3.2	0.003
	0040	2231	223F	SAR001	3.0	0.003
	006D	4300	4300	DMP013	3.1	0.003

3					9.2	
4	003F	2110	2111	MVT01A	2.8	0.002
	004A	2E00	2E04	MVT01B	3.2	0.002
	0045	27B0	27B1	MVT01C	2.9	0.002
	0046	2840	2842	SPLXD4	3.1	0.002

4					11.9	
5	004A	2EA1	2EAD	SPLXD1	3.8	0.001
	0042	24D1	24DF	BSDS01	3.2	0.001

5					7.0	

4-71B Listing of All Exception Volumes, by Cluster

Volume Group Activity by Workload / Chart (STGEHB)

Source: BATWDA file at the DAYS timespan
HARDVA file at the DAYS timespan
WLMSEC file at the DAYS timespan

Function: Provides a graphic view of the I/O activity against each DASD volume group, based on the EXCP counts stored by CA MICS in the BATWDA and HARDVA files.

Each user of the volume group is identified by the WLM Service Class name. This provides the ability to monitor each volume group during various time periods.

Features: For SMS users, grouping for this report is automatically done based on the storage group assigned to each volume.

For non-SMS users, grouping is done by means of a volume table that is created by the StorageMate product and maintained by you.

Even after volume pooling is implemented, both SMS and non-SMS users can use this volume table to simulate proposed changes before they are made.

Three different CA MICS files are combined to produce this chart. The user can set the summarization level to day, zone, or hour.

The input data can also be filtered by hour, zone, data and/or the name of the volume group. This provides a wide variety of data views, from very broad to very narrow.

Both color graphic and printer plots are supported.

SMS
Issues: Prior to SMS implementation, this report could be used to simulate storage groups, making sure the I/O activity against these groups is distributed as expected.

After implementation, the same report could be used to monitor and tune each group, and to simulate proposed changes to the groups.

COLUMN DESCRIPTIONS

Date

The date the activity shown on this chart occurred.

Volume Group Name

The name of the Volume Group supplying the information for this chart. Volume group is determined automatically for SMS users; it is the SMS Storage Group. For non-SMS users, volumes can be assigned to volume pools (at least for reporting purposes) by specifying a Volume Group Table name on the report options screen.

Zone

If summarization by ZONE was specified on the report options screen, the data on this chart will only apply for this ZONE.

Hour

If summarization by Hour was specified on the report options screen, the data on this chart will only apply for this hour.

Pie Slices

Each slice represents the number of I/O requests for this volume group received during this measurement interval that originated from a particular workload. The percentage value represents the percentage of all requests received that originated from each workload.

Pie Labels

Each slice identifies the name of the workload that issued I/O requests to the volumes in this volume group. Similar tasks in MVS are assigned to a common WLM Service Class, which is controlled and reported by means of a service class name. A value of '(SYSTEM)' in the label is used to represent uncaptured I/O requests. This refers to I/O requests that are captured in the totals from RMF for the device, but are not captured by SMF and assigned to the proper user. This is a 'normal' occurrence for all MVS systems, as some work performed by the operating system is considered to be overhead, and is not captured in job accounting.

Figure 4-72. Volume Group Activity by Workload / Chart

Volume Group Allocation by User / Chart (STGEHC)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
VCAVOA file at the DETAIL timespan

Function: Allows users to obtain a visual summary of how each DASD volume group is allocated. This summarization can be done using one or more of the account fields, one or more nodes of the data set name, or using a user-defined table that translates the first node of the data set name into an owner description. The plot makes it easy to identify the major users of each volume group.

Features: For SMS users, grouping for this report is automatically done based on the storage group assigned to each volume.

For non-SMS users, grouping is done by means of a volume table that is created by the StorageMate product and maintained by you.

Even after volume pooling is implemented, both SMS and non-SMS users can use this volume table to simulate proposed changes before they are made.

All current VCA files are combined for this chart. The VCADAA file provides non-VSAM data set information, the VCA_VS file provides VSAM data set information, and the VCAVOA file provides free space information for each volume within the group. The ability to select the summarization element allows users to have many different views of the same data.

This chart detects potential allocation problems at the volume group level and identifies users that are allocating data sets on the wrong groups.

Both color graphic and printer plots are supported.

SMS

Issues: Prior to SMS implementation, this chart can be used to simulate storage groups, making sure that volumes with similar allocation patterns are grouped together.

After implementation, the same chart can be used to monitor and tune each group, and to simulate proposed changes to the groups.

COLUMN DESCRIPTIONS

Volume Group Name

The volume group to which the information on this chart will apply. Volume group is determined automatically for SMS users, because it is the same as the SMS storage group. For non-SMS users, volumes can be assigned to volume pools (at least for reporting purposes) by specifying a Volume Group Table name on the report options screen.

Pie Slices

The pie slices represent the total amount of space in the volume pool allocated to each user, expressed in millions of bytes (megabytes), and also as a percentage of the total.

Pie Labels

The labels on each slice identify the user of the volume group. Based on values coded on the report options panel, this identification can be built using one or more nodes of the data set name, one or more of the VCA account fields, or using an external Data Set Name Table to associate a data set with its owner. A value of '====> FREE SPACE ' is used to represent the portion of the volume pool that is not assigned to any user, and is available for new allocations.




Figure 4-74. Volume Group Allocation by User / Chart

Volume Group Allocation by User / Summary (STGEHC)

Source: VCADAA file at the DETAIL timespan
VCA_VS file at the DETAIL timespan
VCAVOA file at the DETAIL timespan

Function: Concisely summarizes how each DASD volume group is allocated. This summarization can be done using one or more of the account fields, one or more nodes of the data set name, or using a user-defined table that translates the first node of the data set name into an owner description.

Features: For SMS users, grouping for this report is automatically done based on the storage group assigned to each volume.

For non-SMS users, grouping is done by means of a volume table that is created by the StorageMate product and maintained by you.

Even after volume pooling is implemented, both SMS and non-SMS users can use this volume table to simulate proposed changes before they are made.

All current VCA files are combined for this report. The VCADAA file provides non-VSAM data set information, the VCA_VS file provides VSAM data set information, and the VCAVOA file provides free space information for each volume within the group. The ability to select the summarization element allows users to have many different views of the same data.

The report lists all users who have data sets allocated on each volume group, with the users occupying the most space listed first.

The report also detects potential allocation problems at the volume group level and identifies users that are allocating data sets on the wrong groups.

Both color graphic and printer plots are supported.

Generation of the summary report is automatic.

SMS

Issues: Prior to SMS implementation, this chart can be used to simulate storage groups, making sure that volumes with similar allocation patterns are grouped together.

After implementation, the same chart can be used to monitor and tune each group, and to simulate proposed changes to the groups.

COLUMN DESCRIPTIONS

Volume Group Name

The volume group to which the information that follows will apply. Volume group is determined automatically for SMS users, as it is the same as the SMS storage group. For non-SMS users, volumes can be assigned to volume pools (at least for reporting purposes) by specifying a Volume Group Table name on the report options screen.

User

The identification of this user of the volume group. Based on values coded on the report options panel, this identification can be built using one or more nodes of the data set name, one or more of the VCA account fields, or using an external Data Set Name Table to associate a data set with its owner. A value of '====> FREE SPACE <====' under this column is used to represent unallocated space in the volume group.

Percent of all Space

The percentage of this volume group that is allocated to data sets belonging to this user. For the Free Space entry, this represents the percentage of the volume group that is unused. Entries under this column are sorted so that the primary users of the volume group appear first.

Total Space (MB)

The total amount of space, expressed in millions of bytes (megabytes) allocated from this volume group on behalf of this user. For the Free Space entry, this represents the amount of space in the volume group that is unused. Entries under this column are sorted so that the primary users of the volume group appear first.

Percent of all Data Sets

The percentage of all the data sets allocated on this volume group that belong to this user. This will contain a value of zeros for the Free Space entry.

Total Number of Data Sets

The number of data sets allocated on this volume group that belong to this user. This will contain a value of zeros for the Free Space entry.

Minimum Size (KB) Data Set

The smallest data set allocated on this volume group that belongs to this user, expressed in thousands of bytes (kilobytes). This will contain a value of zeros for the Free Space entry. This will also contain zeros for users that have data sets with VTOC entries but no tracks allocated.

Maximum Size (KB) Data Set

The largest data set allocated on this volume group that belongs to this user, expressed in thousands of bytes (kilobytes). This will contain a value of zeros for the Free Space entry.

VOLUME GROUP ALLOCATION BY USER (STGEHC)							1
CA MICS StorageMate							
ABC CORPORATION							
----- System Identifier=TST1 Volume Group Name=SGBYMN -----							
User	Percent of all Alloc Space	Total Space (MB) Allocated	Percent of all Data Sets	Total Number of Data Sets	Minimum Size (KB) Data Set	Maximum Size (KB) Data Set	
SALES	72.2	1365.8	99.1	223	57	100805	
====> FREE SPACE <====	27.5	520.2	0.0	0	0	0	
SYSTEMS-PROG	0.1	2.5	0.9	2	850	1700	
-----		-----		-----			
STORNAME		1888.6		225			
----- System Identifier=TST1 Volume Group Name=SGBYMY -----							
User	Percent of all Alloc Space	Total Space (MB) Allocated	Percent of all Data Sets	Total Number of Data Sets	Minimum Size (KB) Data Set	Maximum Size (KB) Data Set	
====> FREE SPACE <====	49.4	3733.8	0.0	0	0	0	
PROD-CONTROL	24.5	1851.5	17.8	139	0	664375	
MAINTENANCE	7.7	585.3	12.3	96	47	265813	
SALES	5.0	375.1	19.2	150	0	45055	
FINANCE-PAYROLL	3.9	292.8	17.9	140	0	33469	
MANUFACTURING	3.5	265.5	16.0	125	0	28340	
IMS-TESTING	2.9	217.8	6.0	47	0	94949	
SYSTEMS-PROG	1.9	144.6	9.2	72	0	21363	
CUST-SUPPORT	0.3	25.7	2.3	18	47	11964	
IMS-DEVELOPMENT	0.2	16.2	1.7	13	0	10777	
RESEARCH-AND-DEV	0.2	13.4	2.6	20	0	7121	
MARKETING	0.1	8.2	1.0	8	47	5839	
STORAGE-ADMIN	0.1	4.5	0.3	2	190	4273	
IMS-TESTING-ALT	0.0	0.1	0.1	1	95	95	
DATA-CENTER	0.0	0.0	0.1	1	47	47	
-----		-----		-----			
STORNAME		7534.5		832			

Figure 4-75. Volume Group Allocation by User / Summary

Volume Group Configuration / Totals (STGEHA)

Source: VCAVOA file at the DETAIL timespan
HARDVA file at the DAYS timespan

Function: Concisely summarizes performance and capacity information at the volume group level.

Many installations tend to pool or group volumes with similar attributes and functions. This pooling can be done manually, or with automated software products. Regardless of the method chosen, this report makes it easier to monitor and tune these groups as all reporting is done at the group level rather than at the individual volume level. This makes it easy to monitor the performance and availability of the entire group, or to compare service between groups.

This summary report produces the same totals as the detail report, but omits individual totals for each volume within the group. It is suggested that you use this report for regular monitoring of your groups, and then run the detail report when potential problems are detected.

Users planning for products that do automatic volume pooling can use this report to simulate various groupings in order to find the combination that provides optimum performance and availability.

Features: For SMS users, grouping for this report is automatically done based on the storage group assigned to each volume.

For non-SMS users, grouping is done by means of a volume table that is created by the StorageMate product and maintained by you.

Even after volume pooling is implemented, both SMS and non-SMS users can use this volume table to simulate proposed changes before they are made.

Two different CA MICS files are combined to produce this report; one provides space information and the other provides performance information. The HARDVA file can be filtered at the hour or zone level so that the performance

information reflects a specific time period.

SMS

Issues: Prior to SMS implementation, this report can be used to plan and simulate your SMS storage groups. After implementation, the report is useful in monitoring and tuning these same groups, and to simulate proposed changes to the groups.

COLUMN DESCRIPTIONS

Volume Group Name

The volume group to which the information that follows will apply. Volume group is determined automatically for SMS users, as it is the same as the SMS storage group. For non-SMS users, volumes can be assigned to volume pools (at least for reporting purposes) by specifying a Volume Group Table name on the report options screen.

Number of Volumes

The number of volumes found within the input data that were assigned to this volume group.

Frag. Index

The average volume fragmentation index for all the volumes in this volume group. This is calculated by dividing the sum of all volume fragmentation index values by the number of volumes.

Total Data Sets

The total number of data sets residing in this volume group.

Non-VSAM Data Sets

The number of Non-VSAM data sets residing in this volume group.

VSAM Data Sets

The number of VSAM data sets residing in this volume group.

Total Tracks Allocated

The total number of tracks allocated by all the data sets in this volume group.

Non-VSAM Tracks Allocated

The number of tracks allocated by the Non-VSAM data sets in this volume group.

VSAM Tracks Allocated

The number of tracks allocated by the VSAM data sets in this volume group.

Free Tracks

The total number of unused volume tracks for all volumes within this volume group.

Free Percent

The percentage of tracks within this volume group that are available for the allocation of new data sets.

Average I/O Rate

The total number of I/O requests per second being received by the volumes within this volume group. This value is the sum of the rates for the individual volumes in the group.

Average Response Time

The average time a user of this volume group waits, in thousandths of a second (milliseconds), for an I/O request to complete. This is recalculated at the volume group level, using the same technique used to calculate this value for an individual volume.

VOLUME GROUP CONFIGURATION (STGEHA)												5
CA MICS StorageMate - SUMMARY												
ABC CORPORATION												
----- System Identifier=SYS7 -----												
Volume Group Name	Number of Volumes	Frag. Index	+----- Total	Data Sets Non- VSAM	-----+ VSAM	+----- Total	Tracks Allocated Non- VSAM	-----+ VSAM	+----- Free Tracks	-----+ Per- cent	+--- Average I/O Rate	---+ Resp. Time
(NONE)	9	0.240	493	485	8	56962	56938	24	61392	52.0%	13.11	20.804
CICS-PRD	10	0.830	3427	3383	44	127095	125061	2034	5105	4.0%	4.84	25.681
CICS-TST	2	0.503	128	127	1	23464	23461	3	2964	11.0%	0.21	38.382
CORPORAT	11	0.700	3169	3137	32	136721	132680	4041	8460	6.0%	11.74	18.943
DB2-PRD	4	0.234	756	264	492	49822	22502	27320	3154	6.0%	0.36	18.838
DB2-TST	10	0.484	2399	2300	99	129838	116876	12962	2230	2.0%	6.33	23.556
FINANCE	6	0.547	2052	2018	34	78357	75641	2716	1107	1.0%	3.64	21.701
NETWORK	1	0.275	94	79	15	12424	2907	9517	820	6.0%	0.00	20.326
RESEARCH	5	0.475	1670	1658	12	64847	64582	265	1259	2.0%	10.49	22.204
SUPPORT	5	0.299	964	890	74	57384	37556	19828	8746	13.0%	3.39	19.913
SYSPROG	25	0.256	2020	1771	249	281364	211376	69988	49506	15.0%	17.59	26.182
-----	88	0.441	17172	16112	1060	1018278	869580	148698	144743	12.0%	71.71	22.641

Figure 4-76. Volume Group Configuration / Totals

Volume Group Configuration / Detail (STGEHA)

Source: VCAVOA file at the DETAIL timespan
HARDVA file at the DAYS timespan

Function: Concisely summarizes performance and capacity information at the volume group level. This information is reported for each volume and for the entire volume group.

Many installations tend to group volumes with similar attributes and functions. This pooling can be done manually, or with automated software products. Regardless of the method chosen, this report makes it easier to monitor and run these groups, as it reports at both the individual volume and the group level. This makes it easy to monitor the performance and availability of the entire group, or to detect problems with individual volumes in a group.

Users planning for products that do automatic volume pooling can use this report to simulate various groupings in order to find the combination that provides optimum performance and availability.

Features: For SMS users, grouping for this report is automatically done based on the storage group assigned to each volume.

For non-SMS users, grouping is done by means of a volume table that is created by the StorageMate product and maintained by you.

Even after volume pooling is implemented, both SMS and non-SMS users can use this volume table to simulate proposed changes before they are made.

Two different CA MICS files are combined to produce this report; one provides space information and the other provides performance information. The HARDVA file can be filtered at the hour or zone level so that the performance information reflects a specific time period.

SMS

Issues: Prior to SMS implementation, this report can be used to plan and simulate your SMS storage groups.

After implementation, the report is useful in monitoring and tuning these same groups, and to simulate proposed changes to the groups.

COLUMN DESCRIPTIONS

Volume Group

The volume group to which the information that follows will apply. Volume group is determined automatically for SMS users, as it is the same as the SMS storage group. For non-SMS users, volumes can be assigned to volume pools (at least for reporting purposes) by specifying a Volume Group Table name on the report options screen.

Device Address

The physical address of an I/O device.

C.U. Path

The control unit path (MVS/XA users) or logical channel (non-XA users) through which this volume is attached to the I/O subsystem. For some volume pools, it is suggested that volumes in the pool be connected through several paths, to reduce the impact of a hardware failure. For other volume groups, it may be better to isolate the path connections to a smaller group.

Volume Serial Number

The volume serial number of the storage medium.

Device Type

The type of I/O device.

Frag. Index

A measure of how fragmented the allocated space is on the current volume.

Total Data Sets

The total number of data sets residing on this volume.

Non-VSAM Data Sets

The number of Non-VSAM data sets residing on this volume.

VSAM Data Sets

The number of VSAM data sets residing on this volume.

Total Tracks Allocated

The total number of tracks allocated by all the data sets on this volume.

Non-VSAM Tracks Allocated

The number of tracks allocated by the Non-VSAM data sets on this volume.

VSAM Tracks Allocated

The number of tracks allocated by the VSAM data sets on this volume.

Free Tracks

The amount of space on this volume that is available for allocation by data sets.

Percent Free

The percent of space that is available for allocation on this volume.

Average I/O Rate

The average number of I/O requests per second being received by this volume.

Average Response Time

The average time a user of this volume waits, in thousandths of a second (milliseconds), for an I/O request to complete.

VOLUME GROUP CONFIGURATION (STGEHA)														1
CA MICS StorageMate - DETAIL														
ABC CORPORATION														
----- System Identifier=SYS7 / Volume Group=(NONE) -----														
Device Address	C.U. Path	Volume Serial Number	Device Type	Frag. Index	-----+ Total	Data Sets Non-VSAM	-----+ VSAM	--- Tracks Total	Allocated Non-VSAM	---+ VSAM	--- Free Tracks	---+ Per- cent	+ Average I/O Rate	+ Response Time
140	0001	WORK41	3380	0.314	60	59	1	8069	8066	3	5061	38%	2.42	22.3
251	0002	IPL48B	3380	0.392	127	127	0	12930	12930	0	329	2%	0.13	16.4
252	0002	WORK44	3380	0.284	41	40	1	7334	7331	3	5865	44%	0.62	22.1
255	0002	WORK45	3380	0.196	40	39	1	4020	4017	3	9179	69%	1.00	22.6
25B	0002	WORK46	3380	0.114	37	36	1	3881	3878	3	9093	70%	0.68	23.4
25E	0002	WORK47	3380	0.226	33	32	1	3655	3652	3	9477	72%	0.62	21.1
25F	0002	WORK48	3380	0.145	30	29	1	4886	4883	3	8193	62%	0.73	22.5
322	0003	WORK42	3380	0.222	59	58	1	8350	8347	3	4773	36%	4.29	19.3
5A0	000B	WORK43	3380	0.264	66	65	1	3837	3834	3	9422	71%	2.63	19.9
-----				-----										
009 Volumes				0.240	493	485	8	56962	56938	24	61392	52%	13.11	20.8
----- System Identifier=SYS7 / Volume Group=CICS-PRD -----														
Device Address	C.U. Path	Volume Serial Number	Device Type	Frag. Index	-----+ Total	Data Sets Non-VSAM	-----+ VSAM	--- Tracks Total	Allocated Non-VSAM	---+ VSAM	--- Free Tracks	---+ Per- cent	+ Average I/O Rate	+ Response Time
143	0001	LNR902	3380	0.948	351	348	3	12876	12710	166	338	3%	0.22	35.3
147	0001	LNR904	3380	0.805	248	245	3	12679	12358	321	535	4%	0.81	24.1
14B	0001	LNR906	3380	0.793	202	201	1	12850	12847	3	364	3%	0.36	24.1
14D	0001	LNR907	3380	0.689	198	197	1	12750	12735	15	464	3%	0.56	27.1
14F	0001	LNR908	3380	0.803	345	339	6	12456	12327	129	758	6%	0.32	23.0
241	0002	LNR909	3380	0.652	159	149	10	12521	12034	487	693	5%	0.73	22.7
5A2	000B	LNR903	3380	0.950	514	513	1	13081	13066	15	133	1%	0.43	35.5
5A4	000B	LNR901	3380	0.896	558	555	3	12700	12246	454	544	4%	0.58	27.8
5B2	000B	LNR900	3380	1.086	509	504	5	12970	12785	185	274	2%	0.39	22.5
5B6	000B	LNR905	3380	0.674	343	332	11	12212	11953	259	1002	8%	0.44	20.7
-----				-----										
010 Volumes				0.830	3427	3383	44	127095	125061	2034	5105	4%	4.84	25.7
----- System Identifier=SYS7 / Volume Group=CICS-TST -----														
Device Address	C.U. Path	Volume Serial Number	Device Type	Frag. Index	-----+ Total	Data Sets Non-VSAM	-----+ VSAM	--- Tracks Total	Allocated Non-VSAM	---+ VSAM	--- Free Tracks	---+ Per- cent	+ Average I/O Rate	+ Response Time
248	0002	DAT001	3380	0.521	39	38	1	12745	12742	3	469	4%	0.11	31.8
326	0003	DAT002	3380	0.485	89	89	0	10719	10719	0	2495	19%	0.10	46.0
-----				-----										
002 Volumes				0.503	128	127	1	23464	23461	3	2964	11%	0.21	38.4

Figure 4-77. Volume Group Configuration / Detail

Volume Problems (STGECA)

Source: VCADAA file at the DETAIL timespan
VCAVOA file at the DETAIL timespan

Function: Identifies problems with your DASD volumes that could lead to a loss of performance or availability.

Improves reliability by revealing unidentified problems before any loss of performance or availability is experienced.

Helps identify potential space shortage or fragmentation problems before they become serious.

Identifies the seek distances between the VTOC and the VVDS and between the VTOC and the index. These are used to identify seek delays that may occur if the distances are too great.

Features: The error flags on the right side of the report identify potential problems quickly, while the information on the left side will give more detailed information. A correction jobstream can also be generated to correct some problems.

SMS

Issues: Users planning for SMS or similar products can correct problems that would delay their implementation, such as potential VVDS size problems.

This report calculates actual and estimated VVDS sizes, based on the number of data sets in the volume. The size estimate is based on the larger record sizes needed if SMS controls the volume.

COLUMN DESCRIPTIONS

Volume Serial Number

The volume serial number of the storage medium.

% Free Space in VTOC

The percentage of unused Data Set Control Blocks (DSCBs) in the Volume Table of Contents (VTOC) for this volume. Each unused DSCB represents one new data set that can be allocated on this volume, assuming free space on the volume itself is

available. If this number falls below 10%, there is a possibility that the VTOC will become full, and no new allocations can be made on the volume.

% Free Space on Volume

The percentage of total tracks on the volume that remain available for the allocation of new data sets. If this number falls below 10%, there is a possibility that the volume will become full, and no new allocations can be made on it.

Frag. Index

A measure of how fragmented the allocated space is on this volume.

VTOC Size Trks

The total number of tracks allocated to the VTOC portion of each volume. The VTOC serves as a table of contents for the volume, and will contain at least one entry for each data set allocated on the volume. The size and location of the VTOC is established when a new volume is initialized before its first use. Failure to specify a large enough value may cause the VTOC to run out of space, so that no new data sets can be allocated on the volume, even if unallocated space is still available.

Index #of Ext

The VTOC Index is an optional structure that may be allocated on the volume to increase the speed with which a data set can be located on the volume. It helps improve volume performance, especially for volumes with a large number of data sets, such as TSO volumes. If this column is zero, no VTOC Index data set exists. If this column contains a value greater than one, the Index has expanded into multiple extents across the volume, and poor performance or Index space shortages could result. Note that a VTOC Index is required for volumes that will be managed by SMS.

Index Trks Alloc

The VTOC Index is an optional structure that can be allocated on the volume to increase the speed with which a data set can be located on the volume. It helps improve volume performance, especially for volumes with a large number of data sets, such as TSO volumes. This column represents

the total number of tracks that are currently being used to contain the index for this volume. If this column is zero, no VTOC Index data set exists. Note that a VTOC Index is required for volumes that will be managed by SMS.

VVDS #of Ext

This column represents the number of track extents (Sections) on this volume used to contain the VSAM Volume Data Set (VVDS). Prior to SMS, a VVDS was required if the volume contained VSAM data sets, as it was used to store additional VSAM information not kept in the catalog. All volumes that will be managed by SMS require a VVDS, even if the volume will contain no VSAM data sets. If this column is zero, no VVDS data set exists. If this column contains a value greater than one, the VVDS has expanded into multiple extents across the volume, and poor performance or VVDS space shortages could result.

VVDS Trks Alloc

This column represents the number of tracks on this volume used to contain the VSAM Volume Data Set (VVDS). Prior to SMS, a VVDS was required if the volume contained VSAM data sets, as it was used to store additional VSAM information not kept in the catalog. All volumes that will be managed by SMS require a VVDS, even if the volume will contain no VSAM data sets. If this column is zero, no VVDS data set exists. If this column contains a small value, it is possible the VVDS could run out of space when new data sets are created on it, or when the volume is converted to SMS management.

VVDS Trks Req'd

This column represents an estimate of the number of tracks that should exist in the VVDS to represent the data sets currently contained on the volume. The calculation used to derive this value uses the larger values that are required if the volume is to be managed by SMS. If no VVDS exists on the volume, you can use this as an approximation if you decide to add one. If the volume already has a VVDS and you are considering SMS management of it, compare this column with the VVDS Trks Alloc column to see if a size problem might exist after conversion to SMS.

Seek VVDS Cyls

The approximate distance, in cylinders, between the start of the VTOC on this volume and the start of the VVDS. Placement of the VTOC and the VVDS in the same general location on the volume can improve performance, as it reduces the distance the read head must travel between the two structures. Large values in this column could indicate potential performance problems, caused by excessive seeking. This value will be zero if no VVDS exists on the volume.

Seek Index Cyls

The approximate distance, in cylinders, between the start of the VTOC on this volume and the start of the VTOC Index. Placement of the VTOC and the Index in the same general location on the volume can improve performance, as it reduces the distance the read head must travel between the two structures. Large values in this column could indicate potential performance problems, caused by excessive seeking. This value will be zero if no VTOC Index exists on the volume.

DIRF Set

This data set is updated when a non-MVS operating system has added or deleted a data set on the volume without updating the free space information.

No INDX

An asterisk (*) will appear under this column if no VTOC Index data set exists for this volume.

INDX Off

An asterisk (*) will appear under this column if a VTOC Index is not active for this volume. This can be caused by one of three conditions:

- 1) No index data set has been allocated on the volume. The No INDX column should contain an asterisk also.
- 2) An index data set has been defined, but it has never been activated using the IBM DSF utility program.
- 3) An index data set was defined and activated, but a hardware or software error has deactivated it. The IBM DSF utility program

should be used to reactivate it.

Alt Trks

An asterisk (*) will appear under this column if one or more of the alternate tracks on the volume have been assigned. Alternate tracks are assigned when a hardware error has occurred, and the defective track is replaced by one of the alternates. Element VOANOATK in the VCAVOA file can be examined to determine the actual number of alternate tracks that have been assigned.

Mis Trks

An asterisk (*) will appear under this column if the VCC collector detected missing tracks when the data for this volume was collected. This can be caused by the volume not being serialized when the data was collected (refer to your VCC run options), or by actual structure errors in the VTOC. Element VOASPACM in the VCAVOA file can be examined to determine the actual number of missing tracks that were detected.

VVDS Size

An asterisk (*) will appear under this column if no VVDS was found on the volume, or if the current VVDS size (VDS Trks Alloc) is smaller than the estimated VVDS size required (VVDS Trks Req'd).

Seek Dist

An asterisk (*) will appear under this column if the sum of the two seek values (Seek VVDS Cyls and Seek Index Cyls) exceeds 50 cylinders. This might indicate potential performance problems due to the locations of the VTOC, VTOC Index and VVDS structures.

VOLUME PROBLEMS (STGECA)																			1
CA MICS StorageMate																			
ABC CORPORATION																			
----- System Identifier=TST1 Device Type=3380-K -----																			
Volume Serial Number	% Free Space in VTOC	% Free Space on Volume	Frag. Index	VTOC Size Trks	+-Index-+ #of Trks Ext Alloc	+----+ #of Trks Ext Alloc	VVDS Trks Alloc	----+ Trks Reqd	+--Seek--+ VVDS Index Cyls Cyls	+-----+ DIRF Set	No INDX	Error INDX Off	Indicators Alt Trks	Mis Trks	-----+ VVDS Size	Seek Dist			
DMSC56	100	99	0.000	59	00 0	00	0	1	0 0		*	*			*				
DMS995	100	99	0.001	29	01 175	01	10	1	10 13										
EIT01A	83	14	0.019	12	01 2	01	10	1	5 1										
EIT02A	94	83	0.051	12	01 2	01	10	1	16 1										
EITS01	74	35	0.039	15	00 0	01	30	1	543 0		*	*				*			
EITS02	54	5	0.289	14	00 0	01	3	2	362 0		*	*				*			
EITS03	67	30	0.039	15	00 0	01	30	1	688 0		*	*				*			
EITS04	72	37	0.046	15	00 0	01	10	1	1766 0		*	*				*			
EITS05	77	59	0.101	15	00 0	01	10	1	325 0		*	*				*			
NCPD81	96	88	0.064	29	00 0	01	10	1	201 0		*	*				*			
SMSCAT	98	46	0.000	14	01 14	01	3	1	8 1			*							
SMSE02	96	42	0.250	44	01 14	01	10	1	14 3			*							
SMSPG1	96	30	0.050	14	01 14	01	3	1	3 1			*							
V10001	87	35	0.258	44	01 9	01	30	2	4 3										
V10003	90	43	0.280	44	01 9	01	30	2	4 3			*							
V10004	88	21	0.328	44	01 9	01	30	2	4 3										
HSM902	51	38	0.402	59	01 15	00	0	5	0 4						*				
SMSJS1	94	0	0.000	1	00 0	00	0	1	0 0		*	*			*				

Figure 4-78. Volume Problems

VSAM Attributes / Summary (STGEFC)

Source: VCA_VS file at the DETAIL timespan

Function: Allows you to identify common patterns of data set attributes within your VSAM data sets. This is useful in identifying data sets with similar attributes that may have similar processing requirements.

This report also helps to identify groups of data sets that have been allocated with attributes that may cause performance problems or poor space utilization.

Features: You can monitor the use of different attributes by your users, identifying those that could result in performance or space problems.

You can also choose to drop those groups that only represent a small percentage of your data sets from the report.

After running this summary report, three additional reports that expand on the information provided in the summary can be requested.

Generation of the summary report is automatic.

SMS

Issues: SMS allows similar data sets to be managed under one data class; this report helps users identify and define those classes.

You can determine the most common groups of data sets, thus allowing data sets in those groups to be automated quickly.

Minimum, maximum, and average space allocation is also computed and reported for each group, as SMS often allows default space values to be assigned based on the attributes of a data set.

COLUMN DESCRIPTIONS

Report Group

Represents a group of data sets that have identical attributes, as defined by the columns 'VSAM Org. Type' through 'Shr Opt Sys' on each

report line. Report Group is important, as it is used in all other VSAM Attributes reports to refer to this common group of data sets. The number of report groups is determined by the characteristics of the data analyzed, as well as by the percentage cutoff values entered on the report option screen.

of Data Sets

The number of data sets within the data analyzed that contained attributes identical to those defined by this group.

% of Data Sets

The percentage of data sets within the data analyzed that were assigned to this report group. Note that any data sets excluded by filtering values on the report option panel are not included when calculating this percentage.

VSAM Org. Type

Defines the VSAM organization type of the data sets in this group. One of the following values should be present:

KS - Key Sequence Data Set
ES - Entry Sequence Data Set
RR - Relative Record Data Set
LS - Linear Data Set

Note that all the data sets assigned to this report group have the same VSAM organization listed under this column.

C. I. Size

Note that all the data sets assigned to this report group have the same control interval size listed under this column.

LRECL

The maximum length of each logical record in this data set. Note that all the data sets assigned to this report group have the same logical record length listed under this column.

Imbed

Note that all the data sets assigned to this report group have the same imbed attribute listed under this column.

Replicate

Contains a value of 'Y' if the replication attribute was specified for this data set, and a value of 'N' if it was not. If this column is blank, the replication attribute is not valid for this type of data set. Note that all the data sets assigned to this report group have the same replicate attribute listed under this column.

Rel. Key Pos.

The relative position of the first byte of the key within the logical record for ISAM files. Note that all the data sets assigned to this report group have the same relative key position listed under this column.

Key Length

The length of the key associated with each data block in a keyed data set. Note that all the data sets assigned to this report group have the same key length listed under this column.

% Free C. I.

Note that all the data sets assigned to this report group have the same free control interval percentage value listed under this column.

% Free C. A.

Note that all the data sets assigned to this report group have the same free control area percentage value listed under this column.

Shr Opt Reg

The cross-region share options for the data sets in this group. This is derived from the first byte of the DAASHROP element. Note that all the data sets assigned to this report group have the same cross-region share option listed under this column.

Shr Opt Sys

The cross-system share options for the data sets in this group. This is derived from the second byte of the DAASHROP element. Note that all the data sets assigned to this report group have the same cross-system share option listed under this column.

Avg Size (KB)

The average size, expressed in thousands of bytes (kilobytes), of all data sets in this reporting group. This is calculated by taking the total amount of tracks allocated by all data sets in this group, and dividing it by the number of data sets. This value may be useful when trying to set a default amount of space that could be assigned to the data sets in this group.

Min Size (KB)

The smallest amount of space allocated, expressed in thousands of bytes (kilobytes), detected for all of the data sets in this reporting group. A value of zero indicates that at least one data set in the group is allocated with zero extents and is, therefore, using no space for that data set. This value may be useful when trying to set a default amount of space that could be assigned to the data sets in this group.

Max Size (KB)

The largest amount of space allocated, expressed in thousands of bytes (kilobytes), detected for all of the data sets in this reporting group. This value may be useful when trying to set a default amount of space that could be assigned to the data sets in this group.

VSAM ATTRIBUTES (STGEFC)																1
CA MICS StorageMate - SUMMARY																
ABC CORPORATION																
Report Group	# of Data Sets	% of Data Sets	VSAM Org. Type	C.I. Size	LRECL	Imbed	Repl-icate	Rel. Key Pos.	Key Length	% Free C.I.	% Free C.A.	Shr Opt Reg	Shr Opt Sys	Avg Size (KB)	Min Size (KB)	Max Size (KB)
1	61	7.6	KS	512	0	N	N	0	4	0	0	2	3	52	47	57
2	61	7.6	KS	4096	7290	N	N	0	4	0	0	2	3	170	47	1709
3	39	4.9	ES	4096	0			0	0	0	0	3	4	1107	142	1700
4	36	4.5	KS	4096	0	N	N	0	76	0	0	2	3	165	47	963
5	33	4.1	KS	1536	0	N	N	0	23	0	0	2	3	4033	712	8500
6	32	4.0	KS	4096	0	N	N	8	4	0	0	2	3	137	47	963
7	26	3.2	KS	512	0	N	N	9	45	0	0	3	4	273	47	1133
8	20	2.5	KS	2048	0	N	N	0	54	0	0	2	3	555	47	5553
9	18	2.2	KS	1024	0	N	Y	9	45	0	0	3	4	660	47	1424

Figure 4-79. VSAM Attributes / Summary

VSAM Attributes / Chart (STGEFC)

Source: VCA_VS file at the DETAIL timespan

Function: Expands on the information in the summary report, identifying the range of space being allocated by the VSAM data sets within an identified group.

Once you have identified groups of data sets with common attributes, it is often useful to determine the amount of space they represent, and the range of that space allocation.

Features: Identification of size patterns within common data set groups can help when calculating current utilization or planning for future growth.

This report can also help identify exceptional data sets that are improperly sized for the function they perform.

Color graphics is supported.

Generation of this report is optional. Only the summary report is produced automatically.

SMS

Issues: Just as the summary report helps identify potential for SMS data classes, this report helps you determine the default amount of space that should be allocated by each of those classes. This will allow your data classes to be established faster, and with less chance of error.

COLUMN DESCRIPTIONS

Report Group

Identifies the reporting group for which this chart is produced. Refer back to the summary report to determine the group of data sets that comprise this report group. A separate chart is produced for each summary group, so that allocation trends for common groups of data sets can be identified.

Vertical Axis

The vertical axis of the chart represents a list of allocation values ranging from 100 kilobytes (thousand bytes) to 50000 kilobytes.

Horizontal Axis

The horizontal axis represents the cumulative number of data sets found within each allocation range. As this is a cumulative value, the chart line should grow longer as it moves down the page, representing almost all the data sets analyzed as it reaches the bottom.

Other

The numeric values to the right of each chart line represent the number of data sets in each range, the cumulative number, the percent of all data sets within each range, and the cumulative percentage.

VSAM Attributes / Node Analysis (STGEFC)

Source: VCA_VS file at the DETAIL timespan

Function: Expands on the information in the summary report, providing a data set name node summary of the VSAM data sets within each group.

This is useful when the user needs to know the name structures of data sets with a certain set of attributes.

Features: Report options allow analysis of up to nine nodes, (1st, 2nd, ...9th), a range of nodes, or the last node of the data set name.

You can specify that only the most commonly occurring nodes should appear on the report.

Once the summary report has identified a group of data sets with similar attributes, this report helps identify those users or groups that own the data.

SMS

Issues: Establishing SMS data classes usually requires two steps: defining the attributes of the class and finding a way to assign each new data set to the correct class. This report helps with the second part of this process, as it identifies patterns that can be used in ACS routines to make class assignments.

COLUMN DESCRIPTIONS

Report Group

Identifies the reporting group for which this analysis report is produced. Refer back to the summary report to determine the group of data sets that comprise this report group. All data sets assigned to this particular report group will have their data set node(s) analyzed by this report.

Data Set Name Pattern (Node)

The data set node(s) represented by this summarization. The contents of this column are determined by the values coded on the report option screen for this report, and by the values found in your data set names for this report

group. The occurrence of leading, trailing, or duplicate periods in this value indicates that the node(s) you specified were not present for the data set names represented by this group. For example, trying to analyze nodes 2-3 of data set name SYS1.PROCLIB would result in a value of 'PROCLIB.', while trying to use nodes 3-5 would result in a value of '..'.

Data Set Count

The number of data set names included in this report group that matched the data set name structure identified in the Data Set Name Pattern column. The sum of this column may not equal the number of data sets in the report group, as very small groups of common node(s) can be dropped from the analysis, as specified on the report options screen.

% of Group Total

The percentage of all the data set names included in this report group that matched the data set name structure identified in the Data Set Name Pattern column. The sum of this column may not equal 100%, as very small groups of common node(s) can be dropped from the analysis, as specified on the report options screen.

Average Size (KB)

The average size of each data set in this report group with this data set name pattern, expressed in thousands of bytes (kilobytes). This is calculated by adding the size of each data set in the group, and dividing by the number of data sets.

Minimum Size (KB)

The size of the smallest data set in this report group with this data set name pattern, expressed in thousands of bytes (kilobytes).

Maximum Size (KB)

The size of the largest data set in this report group with this data set name pattern, expressed in thousands of bytes (kilobytes).

Percentage of Data Sets Less Than:

This represents a cumulative percentage distribution line for all data sets within this

data set name pattern and this report group.
 Values along the line range from 100 kilobytes to
 10 megabytes, and represent the cumulative
 percentage of data sets having sizes within each
 range.

% Other

The percentage of data sets within this data set
 name pattern and this report group that are larger
 than 10 megabytes in size.

VSAM ATTRIBUTES (STGEFC)																	16
CA MICS StorageMate - NODE ANALYSIS																	
ABC CORPORATION																	
----- Report Group (See Summary Report)= 5 -----																	
Data Set Name Pattern (Node)	Data Set Count	% of Group Total	Average Size (KB)	Minimum Size (KB)	Maximum Size (KB)	+-----	100 KB	250 KB	500 KB	750 KB	1 MB	2 MB	4 MB	6 MB	8 MB	10 MB	% Other
PSSS00V	16	48.48	3870	712	8500	0	0	0	0	25	25	44	50	75	75	100	0
PSSS01V	12	36.36	3939	712	8500	0	0	0	0	25	25	33	50	75	75	100	0
NCPSMPE	4	12.12	5595	2833	8500	0	0	0	0	0	0	0	25	50	75	100	0
CIC	1	3.03	1530	1530	1530	0	0	0	0	0	0	100	100	100	100	100	0

Figure 4-81. VSAM Attributes / Node Analysis

VSAM Attributes / Detail (STGEFC)

Source : VCA_VS file at the DETAIL timespan

Function: Expands on the information in the summary report, providing a detailed listing of the VSAM data sets within selected groups.

This is useful when you need detailed information about all the data sets allocated with a certain set of attributes.

Features: Used in conjunction with the summary report, this report identifies individual data sets that are allocated with attributes causing poor performance or poor space utilization.

Because this report can produce a large volume of output, users can specify that detailed information be provided for selected groups, for all groups, or for none of the groups.

SMS

Issues: Establishing SMS data classes usually requires two steps: defining the attributes of the class and finding a way to assign each new data set to the correct class. This report helps with the second part of this process, as it identifies patterns that can be used in ACS routines to make class assignments.

COLUMN DESCRIPTIONS

Report Group

Identifies the reporting group for which this detail report is produced. Refer back to the summary report to determine the group of data sets that comprise this report group. All data sets assigned to this particular report group will be listed on this detail report. Detailed reports of selected report groups can be requested from the report options screen.

Data Set Name

The name of the data set.

SYSID

The translated system identification that was derived from ORGSYSID by the CA MICS SYSID

Translation process.

Volume Serial Number

The volume serial number of the storage medium.

VSAM Org. Type

Defines the VSAM organization type of this data set. One of the following values should be present:

KS - Key Sequence Data Set
ES - Entry Sequence Data Set
RR - Relative Record Data Set
LS - Linear Data Set

Component Type

The type of VSAM component that this data set is.

Create Date

The date that the data set was first allocated to a DASD volume.

Last Used Date

The date that the data set was last referenced.

Space Alloc (KB)

The total amount of space allocated by this data set, expressed in thousands of bytes (kilobytes). This is calculated by multiplying the number of tracks allocated by the track capacity in kilobytes.

Space Used (KB)

The total amount of space used by this data set, expressed in thousands of bytes (kilobytes). This is calculated by multiplying the number of tracks used by the track capacity in kilobytes.

VSAM ATTRIBUTES (STGEFC)								
CA MICS StorageMate - DETAIL								
ABC CORPORATION								
----- Report Group (Refer to Summary Report)=4 -----								
Data Set Name	SYSID	Volume Serial Number	VSAM Org. Type	Comp- onent Type	Create Date	Last Used Date	Space Alloc (KB)	Space Used (KB)
SNS.CNMSGF.INDEX	SYSA	VH1003	KS	INDEX	17JAN90	.	57	0
SNS.CNMSGF.INDEX	SYSA	VH1003	KS	SSET	17JAN90	.	170	0
SNS.BNJLGPR.INDEX	SYSA	VH1004	KS	INDEX	17JAN90	.	57	0
SNS.PDA.BNJLGPR.INDEX	SYSA	VH1004	KS	SSET	17JAN90	.	170	0
SNS.PDA.BNJLGSE.INDEX	SYSA	VS1001	KS	INDEX	17JAN90	.	57	0
SNS.PDA.BNJLGSE.INDEX	SYSA	VS1001	KS	SSET	17JAN90	.	170	0
SNS.TXXA.BNJ36PR.INDEX	SYSA	VS1004	KS	INDEX	17JAN90	.	57	0
SNS.TXXA.BNJ36PR.INDEX	SYSA	VS1004	KS	SSET	17JAN90	.	170	0
SNS.TXXA.BNJ36SE.INDEX	SYSA	VH1003	KS	INDEX	17JAN90	.	57	0
SNS.TXXA.BNJ36SE.INDEX	SYSA	VH1003	KS	SSET	17JAN90	.	170	0
SNS.PROJECT.NPDA.BNJLGPR.INDEX	SYSA	VS1001	KS	INDEX	11OCT89	25JUL90	57	0
SNS.PROJECT.NPDA.BNJLGPR.INDEX	SYSA	VS1001	KS	SSET	11OCT89	25JUL90	227	57
SNS.PROJECT.NPDA.BNJLGSE.INDEX	SYSA	VH1001	KS	INDEX	11OCT89	.	57	0
SNS.PROJECT.NPDA.BNJLGSE.INDEX	SYSA	VH1001	KS	SSET	11OCT89	.	170	0
SNS.PROJECT.SA07.NPDA.BNJLGPR.INDEX	SYSA	VS1002	KS	INDEX	17JAN90	.	57	0
SNS.PROJECT.SA07.NPDA.BNJLGPR.INDEX	SYSA	VS1002	KS	SSET	17JAN90	.	170	0
SNS.PROJECT.SA07.TXXA.BNJ36PR.INDEX	SYSA	VS1002	KS	INDEX	09NOV89	.	57	0
SNS.PROJECT.SA07.TXXA.BNJ36PR.INDEX	SYSA	VS1002	KS	SSET	09NOV89	.	170	0
SNS.PROJECT.SA07.TXXA.BNJ36SE.INDEX	SYSA	VH1001	KS	INDEX	09NOV89	.	57	0
SNS.PROJECT.SA07.TXXA.BNJ36SE.INDEX	SYSA	VH1001	KS	SSET	09NOV89	.	170	0
SNS.PROJECT.TXXA.BNJ36PR.INDEX	SYSA	VS1004	KS	INDEX	11OCT89	25JUL90	57	0
SNS.PROJECT.TXXA.BNJ36PR.INDEX	SYSA	VS1004	KS	SSET	11OCT89	25JUL90	227	57
SNS.PROJECT.TXXA.BNJ36SE.INDEX	SYSA	VH1001	KS	INDEX	11OCT89	.	57	0
SNS.PROJECT.TXXA.BNJ36SE.INDEX	SYSA	VH1001	KS	SSET	11OCT89	.	170	0
VSNS.FILES.NLDM.BNJLGPR.INDEX	SYSA	SMSE02	KS	INDEX	26APR89	03MAY89	47	0
VSNS.FILES.NLDM.BNJLGPR.INDEX	SYSA	SMSE02	KS	SSET	26APR89	03MAY89	142	47
VSNS.FILES.NPDA.BNJLGSE.INDEX	SYSA	SMSE02	KS	INDEX	26APR89	.	47	0
VSNS.FILES.NPDA.BNJLGSE.INDEX	SYSA	SMSE02	KS	SSET	26APR89	.	142	0
VSNS.PROJECT.NPDA.BNJLGPR.INDEX	SYSA	VH1003	KS	INDEX	25SEP89	10OCT89	57	0
VSNS.PROJECT.NPDA.BNJLGPR.INDEX	SYSA	VH1003	KS	SSET	25SEP89	10OCT89	963	0
VSNS.PROJECT.NPDA.BNJLGSE.INDEX	SYSA	VH1002	KS	INDEX	19OCT89	.	57	0
VSNS.PROJECT.NPDA.BNJLGSE.INDEX	SYSA	VH1002	KS	SSET	19OCT89	.	963	0
VSNS.PROJECT.TXXA.BNJ36PR.INDEX	SYSA	VH1004	KS	INDEX	25SEP89	10OCT89	57	0
VSNS.PROJECT.TXXA.BNJ36PR.INDEX	SYSA	VH1004	KS	SSET	25SEP89	10OCT89	283	0
VSNS.PROJECT.TXXA.BNJ36SE.INDEX	SYSA	VH1003	KS	INDEX	25SEP89	10OCT89	57	0
VSNS.PROJECT.TXXA.BNJ36SE.INDEX	SYSA	VH1003	KS	SSET	25SEP89	10OCT89	283	0
-----							-----	-----
GROUP							5933	161
							=====	=====
							5933	161

Figure 4-82. VSAM Attributes / Detail

VSAM Wasted Space / Chart (STGECF)

Source: VCA_VS file at the DETAIL timespan

Function: This facility helps to identify departments or groups that are wasting DASD storage space due to poor VSAM allocation practices. These problems can be caused either by allocating more space than is necessary, or by using other VSAM options incorrectly. The Storage Administrator can perform up to four different types of space analysis with this facility, and can report the results on one combined report or on separate reports. Correction of the problems detected by this facility may recover a large amount of storage space with a small amount of effort.

The term "wasted space" in these reports refers to any space that is allocated (so that it cannot be used by others), and yet is not currently being used to contain data. In some cases, such as the overallocation of a data set to accommodate future growth, this is normal. These reports also focus strictly on wasted space, and do not consider other factors such as performance. Before making any changes to the VSAM data sets appearing on the reports from this facility, it is strongly suggested that other performance and application factors be considered.

Features: Summary reports should be run regularly to monitor wasted space, with more detailed reporting occurring when problems are detected. A wide variety of filtering and reporting options allow many views of the same data to be obtained.

Summary reports can be produced as tabular reports, printer graphics reports or color graphics reports, and can be summarized based on the data set name, the CA MICS accounting codes, or using a Data Set Group table. Detail reports may then be generated to isolate problems to the data set level. If the detail report still does not provide enough information, up to four different analysis reports can be run.

SMS

Issues: Users planning for SMS could identify those groups with the worst space usage values, as they may be

the best candidates for the immediate benefits of SMS. As groups migrate to SMS control, regular use of this report would verify that the expected space savings from SMS are being realized.

ELEMENT DESCRIPTIONS

SYSID

The System Identifier to which this data applies.

Type

Up to four different types of space checking may be requested in one report run. The Type label appears at the top of the chart and identifies the type of wasted space represented by this chart. The following values may appear:

TOTAL

Shows the combined wasted space for all of the types of space analysis requested in this run.

Block/Track Size

Shows wasted space caused by the size of the physical blocks being written on each track.

Imbed Overhead

Shows wasted space caused by the use of the IMBED attribute.

Overallocation

Shows wasted space caused by allocating more space than is necessary to contain the data.

Record/CI Size

Shows wasted space caused by record sizes that are not efficient for the control interval size being used.

Pie Slices

The pie slices represent the total amount of space being wasted by the data sets in each summarization group, in millions of bytes (megabytes).

Pie Labels

The labels on each slice identify the summarization key for this chart. Depending on the values specified on the report option screen, this will represent one or more VCA account

elements, one or more nodes of the data set name,
or will contain a value derived from your Data Set
Group table.




Figure 4-83. VSAM Wasted Space / Chart

VSAM Wasted Space / Summary (STGE CF)

Source: VCA_VS file at the DETAIL timespan

Function: This facility helps to identify departments or groups that are wasting DASD storage space due to poor VSAM allocation practices. These problems can be caused either by allocating more space than is necessary, or by using other VSAM options incorrectly. The Storage Administrator can perform up to four different types of space analysis with this facility, and can report the results on one combined report or on separate reports. Correction of the problems detected by this facility may recover a large amount of storage space with a small amount of effort.

The term "wasted space" in these reports refers to any space that is allocated (so that it cannot be used by others), and yet is not currently being used to contain data. In some cases, such as the overallocation of a data set to accommodate future growth, this is normal. These reports also focus strictly on wasted space, and do not consider other factors such as performance. Before making any changes to the VSAM data sets appearing on the reports from this facility, it is strongly suggested that other performance and application factors be considered.

Features: Summary reports should be run regularly to monitor wasted space, with more detailed reporting occurring when problems are detected. A wide variety of filtering and reporting options allow many views of the same data to be obtained.

Summary reports can be produced as tabular reports, printer graphics reports or color graphics reports, and can be summarized based on the data set name, the CA MICS accounting codes, or using a Data Set Group table. Detail reports may then be generated to isolate problems to the data set level. If the detail report still does not provide enough information, up to four different analysis reports can be run.

SMS

Issues: Users planning for SMS could identify those groups with the worst space usage values, as they may be

the best candidates for the immediate benefits of SMS. As groups migrate to SMS control, regular use of this report would verify that the expected space savings from SMS are being realized.

COLUMN DESCRIPTIONS

SYSID

The System Identifier to which this data applies.

Data Owner

Identifies the owner of the VSAM data sets represented in this summary group. Depending on your report options, this value is set based on the data set name, a combination of CA MICS account codes, or using a Data Set Group table to assign ownership.

Total Space in MB

Represents the total space allocated by all the VSAM data sets belonging to this summary group, expressed in megabytes.

Total Waste in MB

Represents the total amount of space currently being wasted by the VSAM data sets belonging to this summary group. Depending on your report options, this may represent the sum of up to four different wasted space calculations.

% of Alloc Waste

The percentage of the VSAM space allocated by the data sets in this summary group that is being wasted. This is calculated using the two previous columns.

Total Data Sets

The total number of VSAM data sets found for this summary group that were included in this analysis.

Record/CI Waste in MB

Represents the wasted space caused by record sizes that are not efficient for the control interval size being used. This total represents all the data sets in this summary group, and is expressed in megabytes.

% of Owner Waste

The percentage of the total wasted space reported for this summary group that is caused by Record/CI waste. The sum of all four "% of Owner Waste" columns should be 100%.

Imbed Waste in MB

Represents the wasted space caused by the use of the IMBED attribute. This total represents all the data sets in this summary group, and is expressed in megabytes.

% of Owner Waste

The percentage of the total wasted space reported for this summary group that is caused by IMBED waste. The sum of all four "% of Owner Waste" columns should be 100%.

Block/Trk Waste in MB

Represents the wasted space caused by the size of the physical blocks being written on each track. This total represents all the data sets in this summary group, and is expressed in megabytes.

% of Owner Waste

The percentage of the total wasted space reported for this summary group that is caused by block/track waste. The sum of all four "% of Owner Waste" columns should be 100%.

OverAlloc Waste in MB

Represents the wasted space caused by allocating more space than is necessary to contain the data. This total represents all the data sets in this summary group, and is expressed in megabytes.

% of Owner Waste

The percentage of the total wasted space reported for this summary group that is caused by overallocation waste. The sum of all four "% of Owner Waste" columns should be 100%.

VSAM WASTED SPACE (STGECE)													
CA MICS StorageMate													
ABC CORPORATION													
----- SYSID=IMS8 -----													
Data Owner	Total Space in MB	Total Waste in MB	% of Alloc Waste	Total Data Sets	Record/CI Waste in MB	% of Owner Waste	Imbed Waste in MB	% of Owner Waste	Block/Trk Waste in MB	% of Owner Waste	OverAlloc Waste in MB	% of Owner Waste	
SALES	10101	3764	37	2951	123	3	184	5	968	26	2489	66	
MARKETING	5305	2088	39	807	32	2	19	1	431	21	1605	77	
DATA-CENTER	2548	1026	40	592	15	1	11	1	220	21	779	76	
PAYROLL	1904	422	22	287	49	12	3	1	235	56	135	32	
CORPORATE	595	284	48	288	3	1	3	1	45	16	232	82	
FINANCIAL	21	15	71	17	0	0	0	0	0	3	15	97	
REMOTE-OPS	23	11	49	24	0	2	0	0	2	18	9	81	

SYSID	20496	7609		4966	223		220		1902		5263		
	=====	=====		=====	=====		=====		=====		=====		
	20496	7609		4966	223		220		1902		5263		

Figure 4-84. VSAM Wasted Space / Summary

VSAM Wasted Space / Detail (STGECF)

Source: VCA_VS file at the DETAIL timespan

Function: This facility helps to identify departments or groups that are wasting DASD storage space due to poor VSAM allocation practices. These problems can be caused either by allocating more space than is necessary, or by using other VSAM options incorrectly. The Storage Administrator can perform up to four different types of space analysis with this facility, and can report the results on one combined report or on separate reports. Correction of the problems detected by this facility may recover a large amount of storage space with a small amount of effort.

The term "wasted space" in these reports refers to any space that is allocated (so that it cannot be used by others), and yet is not currently being used to contain data. In some cases, such as the overallocation of a data set to accommodate future growth, this is normal. These reports also focus strictly on wasted space, and do not consider other factors such as performance. Before making any changes to the VSAM data sets appearing on the reports from this facility, it is strongly suggested that other performance and application factors be considered.

Features: Summary reports should be run regularly to monitor wasted space, with more detailed reporting occurring when problems are detected. A wide variety of filtering and reporting options allow many views of the same data to be obtained.

Summary reports can be produced as tabular reports, printer graphics reports or color graphics reports, and can be summarized based on the data set name, the CA MICS accounting codes, or using a Data Set Group table. Detail reports may then be generated to isolate problems to the data set level. If the detail report still does not provide enough information, up to four different analysis reports can be run.

SMS

Issues: Users planning for SMS could identify those groups with the worst space usage values, as they may be

the best candidates for the immediate benefits of SMS. As groups migrate to SMS control, regular use of this report would verify that the expected space savings from SMS are being realized.

COLUMN DESCRIPTIONS

SYSID

The System Identifier to which this data applies.

Data Set Name

Indicates the name of the VSAM data set represented by the statistics on this line of the report. For VSAM clusters, this line may identify the individual cluster component (such as Data or Index), rather than the cluster name.

Data Set Size in MB

Represents the total amount of space being allocated by this data set, expressed in megabytes.

TOTAL Waste in MB

Represents the total amount of space currently being wasted this data set. Depending on your report options, this may represent the sum of up to four different wasted space calculations.

% Alloc Space Waste

The percentage of the space allocated by this data sets that is being wasted. This is calculated using the two previous columns.

Record/CI Waste in MB

Represents the wasted space in this data set caused by record sizes that are not efficient for the control interval size being used. This value is expressed in megabytes.

Imbed Waste in MB

Represents the wasted space in this data set caused by the use of the IMBED attribute. This value is expressed in megabytes.

Block/Trk Waste in MB

Represents the wasted space in this data set caused by the size of the physical blocks being written on each track. This value is expressed in megabytes.

OverAlloc Waste in MB

Represents the wasted space in this data set caused by allocating more space than is necessary to contain the data. This value is expressed in megabytes.

VSAM WASTED SPACE (STGECE)							
CA MICS StorageMate - DETAIL							
ABC CORPORATION							
----- SYSID=IMS8 -----							
Data Set Name	Data Set Size in MB	TOTAL Waste in MB	% Alloc Space Wasted	Record/CI Waste in MB	Imbed Waste in MB	Block/Trk Waste in MB	OverAlloc Waste in MB
COS.PROD.DB5.SDDS.DATA	1309	284	22	26	70	171	18
SYS9.FIN.TEST.DATA	427	228	53	1	0	10	217
COS.PROD.DB2.SDDS.DATA	877	183	21	18	47	116	3
SLSTST.CD1TRB.DATA	672	148	22	6	0	80	62
TEST.VS24RS1.LOCAL.DATA	170	147	87	0	0	0	147
COS.PROD.DB0.SDDS.DATA	409	130	32	6	18	45	61
SARAMEGA.DSNDBD.DSN2SPS.TEST	150	129	86	0	0	0	128
SYS9.PAY.TEST.DATA	214	121	57	1	0	8	112
DB2DSNP.DSNDBD.TESTINIT.ONE	133	115	87	0	0	0	115
INFOTST1.VSAM.DATA	332	103	31	31	0	39	33
COS.TEST.CONVERT.SDDS.DATA	238	85	36	4	10	24	47
COS.PROD.DB0.SDIDS.DATA	119	84	71	0	2	3	79
SYS5.TST33.BANK.DATA	102	78	76	0	0	2	76
TEST.T91C3420.VDD86261.DETAIL	76	66	86	0	0	0	66
SYS6.DSNDBD.RELOAD.DATA	340	63	18	0	0	42	21

SYSID	5567	1965		93	146	540	1186
	=====	=====		=====	=====	=====	=====
	5567	1965		93	146	540	1186

N=	15						

Figure 4-85. VSAM Wasted Space / Detail

VSAM Wasted Space / Analysis (STGECF)

Source: VCA_VS file at the DETAIL timespan

Function: This facility helps to identify departments or groups that are wasting DASD storage space due to poor VSAM allocation practices. These problems can be caused either by allocating more space than is necessary, or by using other VSAM options incorrectly. The Storage Administrator can perform up to four different types of space analysis with this facility, and can report the results on one combined report or on separate reports. Correction of the problems detected by this facility may recover a large amount of storage space with a small amount of effort.

The term "wasted space" in these reports refers to any space that is allocated (so that it cannot be used by others), and yet is not currently being used to contain data. In some cases, such as the overallocation of a data set to accommodate future growth, this is normal. These reports also focus strictly on wasted space, and do not consider other factors such as performance. Before making any changes to the VSAM data sets appearing on the reports from this facility, it is strongly suggested that other performance and application factors be considered.

Features: Summary reports should be run regularly to monitor wasted space, with more detailed reporting occurring when problems are detected. A wide variety of filtering and reporting options allow many views of the same data to be obtained.

Summary reports can be produced as tabular reports, printer graphics reports or color graphics reports, and can be summarized based on the data set name, the CA MICS accounting codes, or using a Data Set Group table. Detail reports may then be generated to isolate problems to the data set level. If the detail report still does not provide enough information, up to four different analysis reports can be run.

SMS

Issues: Users planning for SMS could identify those groups with the worst space usage values, as they may be

the best candidates for the immediate benefits of SMS. As groups migrate to SMS control, regular use of this report would verify that the expected space savings from SMS are being realized.

COLUMN DESCRIPTIONS

Report Types

There are four different types of analysis reports, corresponding with the four types of wasted space checking done by this facility. Generation of the four types of analysis reports may be done in any combination, using options on the report options screen. It is also possible to have analysis done without an analysis report being produced. The four types of analysis reports are as follows:

RECORD/CI ANALYSIS

When defining a VSAM data set, the user specifies, or allows to default, both the record length and the control interval size. If the record size does not fit evenly into the control interval, some space will be wasted within each control interval. This wasted space may be reclaimed by selecting better values for these parameters.

IMBED ANALYSIS

While use of the IMBED VSAM option is often done for performance reasons, it wastes space by reserving one track from every control area for the embedded keys. This space could be reclaimed by using a larger control area, although that may be difficult as the user has no direct control over this. A better approach may be to examine those large data sets that use embedded keys to see if an alternate method would be possible.

BLOCK/TRACK ANALYSIS

Based on allocation parameters, VSAM calculates the size of the physical block that will be written for the data set. Space may be wasted if a block size is selected that does not fit well on the device being used. In general, the larger the block size, the more efficient the use of the track. Some of this space may be reclaimed by specifying

larger values for some options, such as control interval size, that will cause VSAM to use a larger physical block size.

OVERALLOCATION ANALYSIS

As with an data set, it is possible to allocate a VSAM data set with more space than is needed to hold the data. Some overallocation is not bad, as this gives the data set room to grow as more data is added. But data sets with much overallocation should be investigated, especially if they are large. Some of this space may be reclaimed by reallocation the data set with requesting a smaller amount of storage space.

While the example shows only one of the four types of analysis reports, all four types will be described here. As the left side of each report is similar, it will be described in a common section. The unique columns on each of the four types of reports will then be described.

COMMON DATA COLUMNS

The following data columns appear on the left side of each report, and are common to each of the four analysis reports:

Data Set Name

Indicates the name of the VSAM data set represented by the statistics on this line of the report. For VSAM clusters, this line may identify the individual cluster component (such as Data or Index), rather than the cluster name.

Data Set Size in MB

Represents the total amount of space being allocated by this data set, expressed in megabytes.

TOTAL Waste in MB

Represents the total amount of space currently being wasted this data set. This value may represent up to four different types of wasted space.

% Alloc Space Waste

The percentage of the space allocated by this data sets that is being wasted. This is calculated using the two previous columns.

RECORD/CI ANALYSIS

The following data columns are unique to the Record/CI Analysis report. They appear on the right side of the report, after the common columns:

Record/CI Waste in MB

Represents the wasted space in this data set caused by record sizes that are not efficient for the control interval size being used. This value is expressed in megabytes.

Control Interval Size

Lists the control interval size for this data set, expressed in bytes.

Better CI Size

If possible, the report will attempt to calculate a better control interval size than the one being used. If this column is not blank, you may wish to try the suggested value. The report only attempts to calculate this using the most common control interval sizes (512K, 1024K, 2048K, 4096K and 8192K), so you be able to calculate your own better value based on the requirements of the application that uses the data.

Record Size

Indicates the logical record length (LRECL) specified when the data set was created.

% of each CI Wasted

Indicates the percentage of each control interval that is wasted, due to the size of the records being used.

IMBED ANALYSIS

The following data columns are unique to the Imbed Analysis report. They appear on the right side of the report, after the common columns:

Imbed Waste in MB

Represents the wasted space in this data set caused by the use of the IMBED attribute. This value is expressed in megabytes.

Data Set Format

Indicates the VSAM data set format for this data set.

Control Area Size in Bytes

Indicates the VSAM control area size, expressed in bytes.

Control Area Size in Tracks

Indicates the VSAM control area size, expressed in tracks.

% of each CA Wasted

Indicates the percentage of each control area that is being used for embedded keys and cannot be used for other data storage.

BLOCK/TRACK ANALYSIS

The following data columns are unique to the Block/Track Analysis report. They appear on the right side of the report, after the common columns:

Block/Trk Waste in MB

Represents the wasted space in this data set caused by the size of the physical blocks being written on each track. This value is expressed in megabytes.

Device Type

Indicates the type of device on which this data set is stored.

Device Track Size

Indicates the maximum amount of data that may be stored on one track for this type of device.

Block Size

Indicates the physical block size used by VSAM for storing the records within this data set.

Blocks Per Track

Indicates the number of blocks that are written on each track for this data set.

% of each Trk Wasted

The percentage of each track that cannot be used for data storage, due to the physical block size and the gaps that exist between physical blocks. In general, the smaller the block size, the more waste that will occur.

OVERALLOCATION ANALYSIS

The following data columns are unique to the Overalllocation Analysis report. They appear on the right side of the report, after the common columns:

OverAlloc Waste in MB

Represents the wasted space in this data set caused by allocating more space than is necessary to contain the data. This value is expressed in megabytes.

Data Set Format

Indicates the VSAM data set format for this data set.

Allocated Space in Tracks

Indicates the total amount of space allocated for this data set, expressed in DASD tracks.

Used Space in Tracks

Indicates the amount of space allocated to this data set that is actually being used for the storage of data, also expressed in DASD tracks.

Percent Space Unused

Indicates the percentage of the allocated space that is not being used for data storage. This value is calculated from the previous two columns.

VSAM WASTED SPACE (STGE CF)								
CA MICS StorageMate - RECORD/CI ANALYSIS								
ABC CORPORATION								
----- SYSID=IMS8 -----								
Data Set Name	Data Set Size in MB	TOTAL Waste in MB	% Alloc Space Wasted	Record/CI Waste in MB	Control Interval Size	Better CI Size	Record Size	% of each CI Wasted
INFOTEST.VSAM.DATA	332	103	31	31	4096	8192	600	12
COS.PROD.DB5.SDDS.DATA	1309	284	22	26	8192		2000	2
COS.PROD.DB2.SDDS.DATA	877	183	21	18	8192		4000	2
INFOTST.ONE.VSAM.DATA	120	30	25	12	4096	8192	600	12
SYS5.AAA.STAINST.SDDS.DATA	68	19	28	8	8192		3500	15
SLSTST.CD1TRB.DATA	672	148	22	6	24576		4856	1
COS.PROD.DB0.SDDS.DATA	409	130	32	6	8192		669	2
DSICTST.DDIRD	94	18	19	5	4096	8192	384	6
SYS9.FIN.ONLI.DATA	142	36	25	5	22528		1800	4
TESTFIN.AMX101B.DATA	107	19	17	4	8192		600	5
TESTFIN.AMX100A.DATA	107	19	17	4	8192		600	5
PRODDEV.TESTING.MCF.STAGE2.DATA	127	61	48	4	8192		640	6
COS.TEST.CONVERT.SDDS.DATA	238	85	36	4	8192		2000	2
TESTFIN.AMX101A.DATA	85	15	17	4	8192		600	5
BTEST.RUN40.DXVBNHG0.DATA	98	22	22	3	32768		2250	4
-----	-----	-----	-----	-----	-----	-----	-----	-----
SYSID	4786	1171		141				
	=====	=====		=====				
	4786	1171		141				
N= 15								

Figure 4-86. VSAM Wasted Space / Analysis

FACILITY PROCEDURES

The remainder of this chapter outlines some special instructions and procedures for using certain of the CA MICS StorageMate facilities. These instructions relate directly to the use of some of the reports described in the first part of this section.

INSTRUCTIONS FOR CREATING AND MAINTAINING THE PATTERN FILE

One facility in CA MICS StorageMate allows you to create and test various data set name patterns to be used in writing a DFSMS ACS (Automatic Class Selection) Routine. If you plan to use this facility, you must create a pattern file in your Work Data Set that will be used as input by the facility.

Multiple pattern files can exist in your Work Data Set. After its creation, the pattern file will probably continue to be updated and refined as you run reports and add new patterns to the file. The creation and maintenance of this file is done external to CA MICS StorageMate, using an edit facility such as ISPF Edit. You will create and maintain these files using member name(s) of your own choosing. The name you select must then be entered on the execution options screen when you run the ACS Routine Definition report.

The syntax rules for the pattern file are as follows:

GENERAL SYNTAX RULES

Positions 1-72 of each input statement are scanned, with positions 73-80 being ignored. We suggest you use these last 8 positions for sequence numbers.

COMMENT STATEMENTS

Comment statements are allowed and must have an asterisk ('*') as the first non-blank character, followed by at least ONE BLANK.

THE PATTERN STATEMENT

The pattern statement contains up to three different items: the pattern (required), the construct name (optional), and the print switch (optional). These items can start in any position, but must be separated by spaces, must appear in the order listed above, and cannot contain any embedded spaces.

DATA SET NAME PATTERNS

The first value on each pattern statement consists of the data set name pattern itself. Each pattern will be compared against your data set names with the goal of finding a match and assigning a construct name to the data set. Patterns are tested in the order they are defined. Each pattern name can be a maximum of 44 characters in length. Patterns greater than eight characters must consist of qualifiers, separated by periods (.), with not more than eight characters per qualifier. Patterns containing none of the special characters described in this paragraph are considered fully-qualified data set names, and must match exactly. Three special characters can be used. A % (percent) represents a single character, and will match any one position. An * (asterisk) within a qualifier represents 0-8 characters, or can also be used to represent a single qualifier. A ** (double asterisk) represents zero or more qualifiers, and cannot appear with any other characters in the same qualifier. More than three adjacent asterisks is considered an error. Refer to the IBM Manual "MVS/ESA Storage Administration Reference" for more detailed explanations of the syntax just described. Refer also to the example later in this section.

THE CONSTRUCT

The pattern can, optionally, be followed by a 1-8 character construct name. This name will be assigned to each data set whose name matches the pattern defined previously on the same statement. The construct name should consist of alphabetic (A-Z) or numeric (0-9) characters, and should not contain special characters or embedded blanks. The construct name also cannot be the values PRINT or NOPRINT, as those are reserved keywords. If no construct name is specified, a value of "(NONE)" will be assigned.

THE PRINT SWITCH

The pattern or construct name can, optionally, be followed by a keyword with two possible values: Specify PRINT to receive a detailed report of each data set name that matches this pattern or NOPRINT for no detailed report for this pattern. NOPRINT is the default if this is left blank.

UNMATCHED DATA SET STATEMENT


A third type of statement that can exist in the pattern file is a keyword used to control the printing of data set names that fail to match any pattern in the file. This keyword, if coded, must exist as a separate statement in the file. Specify UNMATCHED=PRINT to print data sets names that do not match any pattern or UNMATCHED=NOPRINT to suppress printing. No spaces should precede or follow the equal sign in the keyword. The default is to not print unmatched data set names. Specify this value only once in the pattern file. If it occurs more than once, the last option coded will be used.

SAMPLE PATTERN FILE:

```
* SAMPLE PATTERN FILE FOR TESTING
UNMATCHED=PRINT
SYS%.**          SYSPROG NOPRINT
**.*.LIST        WORKDSN NOPRINT
TEMP.**          WORKDSN NOPRINT
PROD%.**.*.MASTER CACHED1 PRINT
IMS.**.*.BMP.PROD CACHED1 PRINT
SYS1.LINKLIB      SPECIAL NOPRINT
%%%.CRIT.**       BACKUP2 PRINT
SYS1.HASP*        SYSPROG NOPRINT
```

PROCEDURES FOR USING THE ACS TEST FACILITY

IBM's ISMF online facility provides a mechanism for testing the operation of your ACS routines. This is a useful facility when you first create your ACS routines, and also for ongoing maintenance. The test facility uses as input a test case library, containing test cases (simulated allocations of new data sets) that will be passed to your ACS routines. The test case library is a Partitioned Data Set (PDS), with each of the test cases existing as a separate member of the library. The output of the test facility is a report file containing the results of the testing for each test case member. Figure 4-74 illustrates the use and data flow of this facility.






Figure 4-87. ACS Test Facility Processing

While this facility is useful, building the test cases and analyzing the output can be time consuming, especially if you plan to use a large number of test cases. Two facilities provided with CA MICS StorageMate can help in this area:

- o The ACS Test Case Creation facility uses CA MICS data to

build a test case library and populate it with test cases.

- o The ACS Testing Results facility reads the output file from the testing facility and summarizes the test results.

You should use the following procedure when using these facilities to test your ACS routines:

- 1) Use the ACS Test Case Creation facility to build test cases from your CA MICS data. The output of the facility is a jobstream in your Work Data Set.
- 2) Submit the jobstream created in the Work Data Set. It will create and populate the test case library. One member will exist in the library for each test case. The member name will have the following structure:

@dddcnnn

where:

ddd = The first 3 characters of the data set
name being tested

c = A letter A-Z

nnn = A numeric value 0-999

A member named ZZZINDEX will also be created that will contain information necessary to cross-reference the test cases with the testing output file.

- 3) From the ISMF main menu, select the option for AUTOMATIC CLASS SELECTION. From the next menu displayed, select the TEST option. Then select TEST again from the next menu.
- 4) On the test facility screen, you must specify:
 - o The name of the CDS library containing the ACS routines to be tested.
 - o The name of the ACS test library. Specify the data set name that was created in Step 2 above.
 - o The names of the numbers to be tested. Specify @*, because this will test all members but the ZZZINDEX member. You can also use this to subset the testing to certain members. For example, you could specify @PAY* to test only test cases for data sets with the prefix 'PAY'.

- o The testing output report data set that will contain the testing results.
- o The combination of the four possible ACS routines to be tested. Testing all four routines in the same pass is possible, but will lengthen the time to complete the test.

When you have invoked the test facility and the test completes, the testing output report data set will be displayed at the terminal. Press END to terminate browse. You will then specify what to do with the testing output report data set. Specify that the data set be kept, but not printed.

- 5) Invoke the ACS Testing Results facility of CA MICS StorageMate. As parameters to this facility, you must supply the name of the test case library created in step 2 above, and also the testing output report data set created in step 4 above. It is important the two data sets specified represent the same data. The specified testing output report data set must represent the testing of the specified test case library. The ACS Test Results facility can now be used to produce a number of reports to evaluate the testing.

Appendix A: ERROR MESSAGES

Messages

The first part of this appendix lists the messages that may appear on the MICSLOG report after a StorageMate report has been executed. These messages are one of three types:

- o Information. This message usually indicates no error, and is usually issued to document information that existed about the environment when the report was run.
- o Warning. A minor error was detected that may cause all or part of the output to be invalid. Examine the error message text and the output reports to see if they are usable.
- o Error. A severe error has occurred that caused processing to terminate immediately. The results produced by the report, if any, will probably not be valid.

```
+-----+  
|S T G 0 0 0 0 1 |  
+-----+
```

```
TEXT: NO DATA IN <file> FILE AFTER DATA FILTERING.  
CHECK INPUT SPECIFICATIONS AND VALUES SPECIFIED  
FOR FILTERING.  
PROCESSING WILL BE ABORTED.
```

```
TYPE: Error
```

```
REASON: After performing data filtering, one of the files  
used to produce this report contained no data, and no  
reports can be produced. In the message text, <file> is  
replaced by the name of the SAS file that is empty.
```

```
ACTION: Processing aborts with a U998 abend. Verify the  
Cycle Number(s) and filtering criteria you have specified.  
Try removing all filtering criteria and specifying a known  
cycle number.
```

```
REFERENCES: None.
```

```
+-----+  
|S T G 0 0 0 0 2 |  
+-----+
```

TEXT: NO DATA IN <file> FILE AFTER DATA FILTERING.
CHECK INPUT SPECIFICATIONS AND VALUES SPECIFIED
FOR FILTERING.
PROCESSING WILL CONTINUE, BUT SOME REPORTS MAY
NOT BE PRODUCED.

TYPE: Warning

REASON: At least one file used to produce output from this
analysis was found to be empty. In the message text, <file>
is replaced by the SAS name of the empty file.

ACTION: Processing continues, but one or more requested
reports may not be produced. Verify the options specified
during report specification to make sure they are correct.

REFERENCES: None.

```
+-----+  
|S T G 0 0 0 0 3 |  
+-----+
```

TEXT: DYNAMIC ALLOCATION FAILURE FOR DATA SET: <dsname>
MEMBER: <member> / ERROR CODE: <error> /
INFO CODE: <info>
VERIFY THAT THE DATA SET NAME AND MEMBER NAME ARE
CORRECT FOR THIS OPERATION.

TYPE: Error

REASON: An attempt to dynamically allocate one of the data
sets needed by this facility was not successful. In the
message text, <dsname> represents the data set name for
which the allocation was attempted, <member> represents the
member name within the data set, <error> and <info>
represent error and information codes from the MVS dynamic
allocation facility.

ACTION: Processing aborts with a U998abend. Make sure the
data set name shown represents an existing cataloged data

set that is not allocated to another user.

REFERENCES: Refer to a list of dynamic allocation error codes in IBM manuals or the ISPF tutorial.

```
+-----+  
|S T G 0 0 0 0 4 |  
+-----+
```

TEXT: TRKCALC FAILURE FOR DATA SET: <dsname>
VOLUME: <volser> / DEV-BLK-KEY: <parm> / RETURN
CODE: <code>.
VERIFY THAT THE DEVICE TYPE, BLOCK SIZE, AND KEY
LENGTH FOR THIS DATA SET ARE VALID.

TYPE: Warning

REASON: An attempt to use the IBM TRKCALC service for this data set failed. In the message text, <dsname> is replaced by the data set name, <volser> by the serial number, <parm> by the parameter list passed to TRKCALC, and <code> by the error code returned.

ACTION: Processing continues, but some of the statistics computed for this data set will not be valid. The first eight characters of the parameter field are the device type, the next five are the block size, and the final three are the key length. Make sure these are valid. Error code 1 represents an error in TRKCALC itself, error code 2 indicates an invalid parameter list, and error code 3 represents an unknown device type. Verify that the device type in the first part of the parameter is correct, and that your operating system supports this type of device.

REFERENCES: None.

```
+-----+  
| S T G 0 0 0 0 5 |  
+-----+
```

TEXT: DATA SET CATALOG STATUS NOT AVAILABLE FOR DATA
SETS ON SYSID: <system>,
VERIFY THAT CATALOG STATUS COLLECTION WAS
REQUESTED WHEN THE DATA WAS GATHERED ON THIS

SYSTEM.
NO CATALOG ERRORS WILL BE DETECTED FOR ANY OF
THESE DATA SETS.

TYPE: Warning

REASON: Catalog status was not collected for all of the
data sets that reside on the indicated system. In the
message text, <system> is replaced by the name of the system
that had no status.

ACTION: Processing continues, but no catalog status errors
will be detected for any of the data sets on the indicated
system. Verify that the control statements used to run
CA MICS data collection on the indicated system specify that
catalog status is to be collected.

REFERENCES: CA MICS DASD Space Collector User Guide.

```
+-----+  
|S T G 0 0 0 0 6 |  
+-----+
```

TEXT: INVALID DATA SET NAME MASK FOUND: <mask> / ERROR
CODE: <code>.
VERIFY THAT THE MASK SHOWN FOLLOWS THE RULES
FOR VALID DATA SET NAME MASKS.
THIS STATEMENT WILL BE IGNORED AND PROCESSING
WILL CONTINUE.

TYPE: Warning

REASON: One of the data set pattern masks in the pattern
file you created is invalid. In the message text, <mask> is
replaced by the pattern mask in error, and <code> is
replaced by an error code of 2 or 3.

ACTION: Processing continues, but the entry shown will be
ignored. If the error code is 2, the pattern violates the
syntax rules for pattern masks, and should be corrected. If
the error code is 3, an internal logic error occurred.

REFERENCES: Online report tutorials. CA MICS StorageMate

Guide, Section 4.3.

```
+-----+  
|S T G 0 0 0 0 7 |  
+-----+
```

TEXT: THE INTERNAL MAXIMUM OF <level> SUMMARIZATION
LEVELS HAS BEEN EXCEEDED FOR THIS REPORT.
MODIFY INPUT, FILTERING, OR SUMMARIZATION
OPTIONS TO CREATE FEWER LEVELS.
PROCESSING WILL CONTINUE, BUT ANY REMAINING
DATA WILL NOT BE INCLUDED IN THE OUTPUT.

TYPE: Warning

REASON: The number of summary levels created has exceeded
the internal limit for the report. In the message text,
<level> is replaced by the number of summary levels
supported.

ACTION: Processing continues, but all summary levels past
the limit shown will be ignored. Modify your report options
to include less data in the analysis or to otherwise reduce
the number of reporting levels.

REFERENCES: None.

```
+-----+  
|S T G 0 0 0 0 8 |  
+-----+
```

TEXT: THE FOLLOWING BLANK ENTRY WAS FOUND WITHIN THE
DATA SET GROUP TABLE:
<entry>
PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: Record positions 1-72 of the Data Set Group Table
entry shown were found to contain all spaces. In the
message text, <entry> is replaced by the image of the entry

that is in error.

ACTION: Table processing will continue, but will issue a U998abend after all statements have been processed. Remove the incorrect statement or change it to a comment by adding an asterisk. You should leave sequence numbers in columns 73-80 of each table entry to help identify these errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+  
|S T G 0 0 0 0 9 |  
+-----+
```

TEXT: THE DSNAME PORTION OF THIS DATA SET GROUP TABLE ENTRY EXCEEDS 44 CHARACTERS:

<entry>

PROCESSING WILL CONTINUE, BUT THIS STATEMENT WILL BE IGNORED.

TYPE: Warning

REASON: The statement shown from your Data Set Group Table is invalid. The data set portion of the statement exceeds the maximum size of 44 characters. In the message text, <entry> is replaced by the image of the entry that is in error.

ACTION: Table processing will continue, but will issue a U998abend after all statements have been processed. Correct the statement that is in error or remove it from the Data Set Group Table. You should leave sequence numbers in columns 73-80 of each table entry to help identify these errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+  
|S T G 0 0 0 1 0 |  
+-----+
```

TEXT: EMBEDDED ASTERISK FOUND WITHIN THE DSNAME PORTION OF THIS DATA SET GROUP TABLE ENTRY:

<entry>
PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: A syntax error was detected in the statement shown from your Data Set Group Table. If the data set name portion of the entry contains an asterisk, it must be the last character. In the message text, <entry> is replaced by the image of the entry that is in error.

ACTION: Table processing will continue, but will issue a U998 abend after all statements have been processed. Correct the statement that is in error or remove it from the Data Set Group Table. You should leave sequence numbers in columns 73-80 of each table entry to help identify these errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+  
|S T G 0 0 0 1 1 |  
+-----+
```

TEXT: GROUP NAME MISSING FROM THE FOLLOWING DATA SET
GROUP TABLE ENTRY:
<entry>
PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: The statement shown from your Data Set Group Table contained a syntax error. No valid group assignment was found as the second operand on the statement. In the message text, <entry> is replaced by the image of the entry that is in error.

ACTION: Table processing will continue, but will issue a U998 abend after all statements have been processed. Correct the statement that is in error or remove it from the Data Set Group Table. You should leave sequence numbers in

columns 73-80 of each table entry to help identify these errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+
|S T G 0 0 0 1 2 |
+-----+
```

TEXT: GROUP NAME FROM THE FOLLOWING DATA SET GROUP TABLE
ENTRY EXCEEDS 40 CHARACTERS:

<entry>

PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: The statement shown from your Data Set Group Table
contained a syntax error. The group assignment specified as
the second operand on the statement exceeds the maximum
limit of 40 characters. In the message text, <entry> is
replaced by the image of the entry that is in error.

ACTION: Table processing will continue, but will issue a
U998 abend after all statements have been processed.
Correct the statement that is in error or remove it from the
Data Set Group Table. You should leave sequence numbers in
columns 73-80 of each table entry to help identify these
errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+
|S T G 0 0 0 1 3 |
+-----+
```

TEXT: NO VALID STATEMENTS WERE FOUND WITHIN THE
SPECIFIED DATA SET GROUP TABLE.

THIS MAY BE CAUSED BY AN EMPTY TABLE OR BY
PREVIOUS ERRORS.

PROCESSING WILL BE ABORTED.

TYPE: Error

REASON: A Data Set Group Table was specified for use by this report, but no valid entries were found within it.

ACTION: Processing Terminates with a U998 ABEND. Refer to the Work Data Set and Data Set Group Table names shown in message STG00024. Verify that the indicated table in the indicated data set contains valid Data Set Group entries. Correct the errors or specify a different table name.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+  
|S T G 0 0 0 1 4 |  
+-----+
```

TEXT: THE FOLLOWING BLANK ENTRY WAS FOUND WITHIN THE VOLUME GROUP TABLE:

<entry>

PROCESSING WILL CONTINUE, BUT THIS STATEMENT WILL BE IGNORED.

TYPE: Warning

REASON: Record positions 1-72 of the Volume Group Table entry shown were found to contain all spaces. In the message text, <entry> is replaced by the image of the entry that is in error.

ACTION: Table processing will continue, but will issue a U998 abend after all statements have been processed. Remove the statement that is in error or change it to a comment by adding an asterisk. You should leave sequence numbers in columns 73-80 of each table entry to help identify these errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
|S T G 0 0 0 1 5 |  
+-----+
```

TEXT: THE VOLUME PORTION OF THIS VOLUME GROUP TABLE
ENTRY EXCEEDS 6 CHARACTERS:

<entry>

PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: The statement shown from your Volume Group Table is
invalid. The volume portion of the statement exceeds the
maximum size of six characters. In the message text,
<entry> is replaced by the image of the entry that is in
error.

ACTION: Table processing will continue, but will issue a
U998 abend after all statements have been processed.
Correct the statement that is in error or remove it from the
Volume Group Table. You should leave sequence numbers in
columns 73-80 of each table entry to help identify these
errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
|S T G 0 0 0 1 6 |  
+-----+
```

TEXT: EMBEDDED ASTERISK FOUND WITHIN THE VOLUME PORTION
OF THIS VOLUME GROUP TABLE ENTRY:

<entry>

PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: A syntax error was detected in the statement shown
from your Volume Group Table. If the volume name portion of
the entry contains an asterisk, it must be the last
character. In the message text, <entry> is replaced by the
image of the entry that is in error.

ACTION: Table processing will continue, but will issue a
U998 abend after all statements have been processed.

Correct the statement that is in error or remove it from the Volume Group Table. You should leave sequence numbers in columns 73-80 of each table entry to help identify these errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
|S T G 0 0 0 1 7 |  
+-----+
```

TEXT: GROUP NAME MISSING FROM THE FOLLOWING VOLUME GROUP
TABLE ENTRY:

<entry>

PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: The statement shown from your Volume Group Table
contained a syntax error. No valid volume group assignment
was found as the second operand on the statement. In the
message text, <entry> is replaced by the image of the entry
that is in error.

ACTION: Table processing will continue, but will issue a
U998 abend after all statements have been processed.
Correct the statement that is in error or remove it from the
Volume Group Table. You should leave sequence numbers in
columns 73-80 of each table entry to help identify these
errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
|S T G 0 0 0 1 8 |  
+-----+
```

TEXT: GROUP NAME FROM THE FOLLOWING VOLUME GROUP TABLE
ENTRY EXCEEDS 8 CHARACTERS:

<entry>

PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: The statement shown from your Volume Group Table contained a syntax error. The volume group assignment specified as the second operand on the statement exceeds the maximum limit of eight characters. In the message text, <entry> is replaced by the image of the entry that is in error.

ACTION: Table processing will continue, but will issue a U998 abend after all statements have been processed. Correct the statement that is in error or remove it from the Volume Group Table. You should leave sequence numbers in columns 73-80 of each table entry to help identify these errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+
|S T G 0 0 0 1 9 |
+-----+
```

TEXT: NO VALID STATEMENTS WERE FOUND WITHIN THE SPECIFIED VOLUME GROUP TABLE.
THIS MAY BE CAUSED BY AN EMPTY TABLE OR BY PREVIOUS ERRORS.
PROCESSING WILL BE ABORTED.

TYPE: Error

REASON: A Volume Group Table was specified for use by this report, but no valid entries were found within it.

ACTION: Processing terminates with a U998 abend. Refer to the Work Data Set and Volume Group Table names shown in a previous message. Verify that the indicated table in the indicated data set contains valid Volume Group entries. Correct the errors or specify a different table name.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+
```



```
|S T G 0 0 0 2 0 |  
+-----+
```

TEXT: DYNAMIC ALLOCATION FAILURE ATTEMPTING TO FREE
DDNAME: <ddname>
ERROR CODE: <error> / INFO CODE: <info>.
PROCESSING WILL CONTINUE, BUT THE DATA SET WILL
REMAIN ALLOCATED.

TYPE: Warning

REASON: An attempt to dynamically free one of the data sets
previously allocated by this facility was not successful.
In the message text, <ddname> represents the allocated
ddname for which the free was attempted, <error> and <info>
represent error and information codes from the MVS dynamic
allocation facility.

ACTION: Processing continues, but the data set allocated to
the ddname shown remains allocated. For a batch job, this
allocation will probably be freed at job termination. TS0
users may want to free the allocation manually using the TS0
FREE command.

REFERENCES: Refer to a list of dynamic allocation error
codes in IBM manuals or the ISPF tutorial.

```
+-----+  
|S T G 0 0 0 2 1 |  
+-----+
```

TEXT: WARNING: THE DATA SET GROUP TABLE JUST CREATED
CONTAINS <count> ENTRIES.
A TABLE THIS SIZE MAY CAUSE PERFORMANCE
PROBLEMS OR SAS MEMORY SIZE ERRORS.
EDIT THE TABLE TO REDUCE THE NUMBER OF ENTRIES
TO A MANAGEABLE NUMBER.

TYPE: Warning

REASON: The Data Set Group Table just created contains a
large number of entries. A large table may cause slow
report performance or may even cause SAS errors or memory
size errors when running reports. In the message text,

<count> is replaced by the number of table entries.

ACTION: You should reduce the table size by creating the table again with different options or by editing the table to combine entries.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+  
|S T G 0 0 0 2 2 |  
+-----+
```

TEXT: A VOLUME GROUP TABLE HAS BEEN CREATED IN MEMBER
<member> OF DATA SET <dsname>.
CREATED BY JOBNAME/USERID <user> ON <date> AT
<time>.
THE TABLE CONTAINS <count> UNIQUE VOLUME
ENTRIES.

TYPE: Information

REASON: This report created a new Volume Group Table in
your Work Data Set. In the message text, <dsname> indicates
the Work Data Set name, <member> represents the table name
created, <user> indicates the TSO user or batch job that
created the table, <date> and <time> indicate when the table
was created, and <count> indicates the number of non-comment
entries in the table.

ACTION: Remember to edit this table (using the CA MICS
StorageMate option facility) to tailor it to your needs
before attempting to use it.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
|S T G 0 0 0 2 3 |  
+-----+
```

TEXT: A DATA SET GROUP TABLE HAS BEEN CREATED IN MEMBER
<member> OF DATA SET <dsname>.
CREATED BY JOBNAME/USERID <user> ON <date> AT

<time>.

THE TABLE CONTAINS <count> UNIQUE DATA SET
GROUP ENTRIES.

TYPE: Information

REASON: This report created a new Data Set Group Table in
your Work Data Set. In the message text, <dsname> indicates
the Work Data Set name, <member> represents the table name
created, <user> indicates the TSO user or batch job that
created the table, <date> and <time> indicate when the table
was created, and <count> indicates the number of non-comment
entries in the table.

ACTION: Remember to edit this table (using the CA MICS
StorageMate option facility) to tailor it to your needs
before attempting to use it.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+  
|S T G 0 0 0 2 4 |  
+-----+
```

TEXT: VOLUME GROUP TABLE <dsname(table)> WAS SPECIFIED
FOR THIS INQUIRY.

TYPE: Information

REASON: As part of the report options for this inquiry a
Volume Group Table was specified, to be used in volume
grouping. In the message text, <table> represents the
specified table name, and <dsname> specifies the Work Data
Set in which the table should reside.

ACTION: None.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
|S T G 0 0 0 2 5 |  
+-----+
```

TEXT: DATA SET GROUP TABLE <dsname(table)> WAS
SPECIFIED FOR THIS INQUIRY.

TYPE: Information

REASON: As part of the report options for this inquiry a
Data Set Group Table was specified, to be used in data set
grouping. In the message text, <table> represents the
specified table name, and <dsname> specifies the Work Data
Set in which the table should reside.

ACTION: None.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+  
|S T G 0 0 0 2 6 |  
+-----+
```

TEXT: THE VOLUME PORTION OF THIS VOLUME GROUP TABLE
ENTRY IS MISSING A VOL= PREFIX:
<entry>
PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: The statement shown from your Volume Group Table in
invalid. The volume portion of the statement does not
contain the required prefix of VOL=. In the message text,
<entry> is replaced by the image of the entry in error.

ACTION: Table processing will continue, but will issue a
U998 abend after all statements have been processed.
Correct the statement in error or remove it from the Volume
Group Table. You should leave sequence numbers in columns
73-80 of each table entry to help identify these errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
|S T G 0 0 0 2 7 |  
+-----+
```

+-----+

TEXT: THE DSNAME PORTION OF THIS DATA SET GROUP TABLE
ENTRY IS MISSING A DSN= PREFIX:

<entry>

PROCESSING WILL CONTINUE, BUT THIS STATEMENT
WILL BE IGNORED.

TYPE: Warning

REASON: The statement shown from your Data Set Group Table
is invalid. The data set portion of the statement does not
contain the required prefix of DSN=. In the message text,
<entry> is replaced by the image of the entry in error.

ACTION: Table processing will continue, but will issue a
U998 abend after all statements have been processed.
Correct the statement in error or remove it from the Data
Set Group Table. You should leave sequence numbers in
columns 73-80 of each table entry to help identify these
errors.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

+-----+

|S T G 0 0 0 2 8 |

+-----+

TEXT: ONE OR MORE ERRORS WERE FOUND IN VOLUME GROUP
TABLE <dsname(table)>.

PROCESSING WILL BE ABORTED.

CORRECT THE ERRORS AND RETRY, OR SPECIFY A
DIFFERENT TABLE.

TYPE: Error

REASON: At least one syntax error occurred when processing
the entries in the Volume Group Table you specified as input
for this report. These errors should have been identified
with previous error messages. In the message text, <table>
is replaced with the table name you indicated, and <dsname>
represents the Work Data Set from which it was read.

ACTION: Processing aborts with a U998abend. Correct the errors identified in previous messages or specify a different table name.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
|S T G 0 0 0 2 9 |  
+-----+
```

TEXT: ONE OR MORE ERRORS WERE FOUND IN DATA SET GROUP
TABLE <dsname(member)>.
PROCESSING WILL BE ABORTED.
CORRECT THE ERRORS AND RETRY, OR SPECIFY A
DIFFERENT TABLE.

TYPE: Error

REASON: At least one syntax error occurred when processing the entries in the Data Set Group Table you specified as input for this report. These errors should have been identified with previous error messages. In the message text, <table> is replaced with the table name you indicated, and <dsname> represents the Work Data Set from which it was read.

ACTION: Processing aborts with a U998abend. Correct the errors identified in previous messages or specify a different table name.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+  
|S T G 0 0 0 3 0 |  
+-----+
```

TEXT: THE FOLLOWING SPACE MANAGEMENT SIMULATION VALUES
WERE SPECIFIED:
<values>

TYPE: Information

REASON: This message indicates the simulation values you specified for this execution of the report. In the message text, <values> is replaced by up to four keywords indicating the simulation values that were specified. Those values for which no value was coded will not appear in the message.

ACTION: None.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
|S T G 0 0 0 3 1 |  
+-----+
```

TEXT: A SAMPLE JOBSTREAM TO CORRECT VOLUME PROBLEMS HAS BEEN CREATED IN THE FOLLOWING DATA SET:
<dsname(member)>
PLEASE REVIEW THIS JOB CAREFULLY
BEFORE SUBMITTING IT FOR EXECUTION.

TYPE: Information

REASON: At your request, a jobstream has been created in your Work Data Set that will correct some of the problems detected by this report. In the message text, <dsname> identifies your Work Data Set, and <member> identifies the name of the member that was created.

ACTION: Review the specified member carefully, then use it to correct some of the problems identified on the report.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
|S T G 0 0 0 3 2 |  
+-----+
```

TEXT: A SAMPLE JOBSTREAM TO CORRECT DATA SET PROBLEMS HAS BEEN CREATED IN THE FOLLOWING DATA SET:
<dsname(member)>
PLEASE REVIEW THIS JOB CAREFULLY
BEFORE SUBMITTING IT FOR EXECUTION.

TYPE: Information

REASON: At your request, a jobstream has been created in your Work Data Set that will correct some of the problems detected by this report. In the message text, <dsname> identifies your Work Data Set, and <member> identifies the name of the member that was created.

ACTION: Review the specified member carefully, then use it to correct some of the problems identified on the report.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
|S T G 0 0 0 3 3 |  
+-----+
```

TEXT: THE FOLLOWING PATTERN FILE WAS SPECIFIED AS INPUT FOR THIS ANALYSIS:
<dsname(member)>

TYPE: Information

REASON: This message documents the name of the pattern file you specified to be used as input for this simulation. In the message text, <dsname> identifies the name of your Work Data Set, and <member> identifies the name of the pattern file that was used.

ACTION: None.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
|S T G 0 0 0 3 4 |  
+-----+
```

TEXT: A SAMPLE ACS ROUTINE HAS BEEN CREATED IN MEMBER <member> OF DATA SET <dsname>.
CREATED BY JOBNAME/USERID <user> ON <date> AT <time>.

TYPE: Information

REASON: At your request, this report created a sample ACS routine, and placed it in your Work Data Set. In the message text, <dsname> indicates the Work Data Set name, <member> represents the ACS routine created, <user> indicates the TSO user or batch job that created the routine, and <date> and <time> indicate when the routine was created.

ACTION: Remember to inspect and modify this routine to tailor it to your needs before attempting to use it as a true ACS routine.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
|S T G 0 0 0 3 5 |  
+-----+
```

TEXT: A LIST OF DASD-TO-TAPE CANDIDATES HAS BEEN CREATED
IN THE FOLLOWING DATA SET:

<dsname(member)>

TYPE: Information

REASON: At your request, this report created a list of data sets that may be good candidates for movement from DASD to tape, and placed it in your Work Data Set. In the message text, <dsname> indicates the Work Data Set name, and <member> represents the member name used to contain the list.

ACTION: Inspect the created list and modify it to meet your needs. Then use it to create jobstreams or user exits that will move the desired data sets to tape.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
|S T G 0 0 0 3 6 |  
+-----+
```

TEXT: THE FOLLOWING EXTERNAL TAPE FILE WAS SPECIFIED AS
INPUT FOR THIS ANALYSIS:

<dsname(TAPEDATA)>

THIS SAS DATA SET SHOULD CONTAIN TAPE-RELATED
DATA IN THE FORMAT DESCRIBED IN THE
DOCUMENTATION.

TYPE: Information

REASON: The inquiry you executed obtained tape-related
information from a non-CA MICS SAS data set that you
specified. This message documents the name of the data set
that you specified using the StorageMate and MICF options
facility. In the message text, <dsname> identifies the name
of the data set that was specified.

ACTION: None.

REFERENCES: CA MICS StorageMate Guide, Appendix C.

```
+-----+  
|S T G 0 0 0 3 7 |  
+-----+
```

TEXT: A LIST OF TAPE-TO-DASD CANDIDATES HAS BEEN
CREATED IN THE FOLLOWING DATA SET:

<dsname(member)>

TYPE: Information

REASON: At your request, this report created a list of data
sets that may be good candidates for movement from tape to
DASD, and placed it in your Work Data Set. In the message
text, <dsname> indicates the Work Data Set name, and
<member> represents the member name used to contain the
list.

ACTION: Inspect the created list and modify it to meet your
needs. Then use it to create jobstreams or user exits that
will move the desired data sets to DASD.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
|S T G 0 0 0 3 8 |  
+-----+
```

TEXT: DATA SET NOT FOUND IN DATA SET GROUP TABLE:
<dsname>

TYPE: Information

REASON: For this inquiry, you specified that a Data Set Group Table should be used for data set grouping. In the process of doing that grouping, a data set was found that did not match any of the entries in the table. In the message text, <dsname> represents the name of the data set that was not found in the table. This message is only produced if you have set the &STGDEBUG option to a value of Y (Yes).

ACTION: The specified data set is assigned a group name of '(UNKNOWN)', and processing continues. Verify that the table you specified (identified in a previous message) was correct. Create a new table, or add entries for the identified data set(s) to your existing table.

REFERENCES: CA MICS StorageMate Guide, Section 3.5.

```
+-----+  
| S T G 0 0 0 3 9 |  
+-----+
```

TEXT: SMSUSE KEYWORD FOUND WITHIN THE VOLUME ENTRIES FOR
THIS VOLUME GROUP TABLE.
IF THIS KEYWORD IS PRESENT, IT MUST PRECEDE ANY
VOLUME ENTRIES.
CORRECT THE TABLE AND RESTART THIS REPORT.

TYPE: Error

REASON: If an SMSUSE keyword appears in the Volume Group Table, it must appear before any of the volume entries.

ACTION: Processing terminates. Modify the Volume Group

Table to place the SMSUSE keyword before any volume entries.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
| S T G 0 0 0 4 0 |  
+-----+
```

TEXT: INVALID SMSUSE KEYWORD FOUND WITHIN THIS VOLUME GROUP TABLE.
VALID OPTIONS ARE SMSUSE=YES OR SMSUSE=NO.
CORRECT THE TABLE AND RESTART THIS REPORT.

TYPE: Error

REASON: An SMSUSE keyword with an invalid option was found within the Volume Group Table you specified. The keyword must be followed by an equal sign, and then the word YES or NO. No spaces should separate the keyword and its assigned value.

ACTION: Processing terminates. Edit the specified Volume Group Table and change the SMSUSE keyword to meet the standards specified above.

REFERENCES: CA MICS StorageMate Guide, Section 3.4.

```
+-----+  
| S T G 0 0 0 4 1 |  
+-----+
```

TEXT: THE MICF FILE FILTERING ROUTINE ENDED WITH AN ERROR CODE OF <errcode>.
MAKE SURE THE FILES NEEDED BY THIS REPORT ARE AVAILABLE IN THE UNIT DATABASES (DBIDS) SPECIFIED.
REPORT EXECUTION IS TERMINATED.

TYPE: Error

REASON: The MICF routine called to drop uninstalled files from the inquiry completed with an unexpected error. In the message text, <errcode> is replaced by the error code

returned from the routine.

ACTION: Processing terminates. See Section 4.3 of this guide to determine the files used by this report, and make sure those files are available in the unit databases (DBIDs) specified.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 4 2 |  
+-----+
```

TEXT: A REQUIRED FILE (NAME = <file>) COULD NOT BE FOUND
IN THE SPECIFIED UNIT DATABASE ID LIST.
MODIFY THE UNIT DATABASE ID LIST TO INCLUDE THE
UNIT WHERE THIS FILE RESIDES.
REPORT EXECUTION IS TERMINATED.

TYPE: Error

REASON: One of the required files for this report could not
be found in the unit database ID list you supplied. In the
message text, <file> identifies the name of the file that
could not be located.

ACTION: Processing terminates. Verify that the required
files for this report are installed in a unit. Change the
database list to include the databases where required
files are located.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 4 3 |  
+-----+
```

TEXT: A REQUIRED FILE WAS NOT FOUND IN ANY UNIT OF
THE SPECIFIED/DEFAULT UNIT DATABASE ID LIST.
DBID LIST: <dbidlst> / RETURN CODE: <retcode>.
REPORT EXECUTION IS TERMINATED.

TYPE: Error

REASON: A required file could not be located within the unit database ID list that you supplied or were provided by MICF as a default. In the message text, <dbidlst> is replaced by the list of DBIDs, and <retcode> is the return code from the MICF validation routine.

ACTION: Processing terminates. Make sure the required files for the report you are executing are installed in a CA MICS unit. Change the DBID specification to include all required units.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 4 4 |  
+-----+
```

TEXT: THE SPECIFIED/DEFAULT UNIT DATABASE ID LIST DOES NOT CONTAIN ALL REQUIRED FILES FOR THIS REPORT.
AN ATTEMPT WILL BE MADE TO CONSTRUCT A NEW UNIT LIST USING ALL VALID UNITS IN THIS COMPLEX.
REPORT EXECUTION CONTINUES.

TYPE: Warning

REASON: The default or supplied unit database ID list for this report did not include units containing all the files needed for this report.

ACTION: MICF will attempt to build a new list from all available units. If that attempt is successful, processing will continue. Otherwise, additional messages will be issued, and processing will fail.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 4 5 |  
+-----+
```

TEXT: UNEXPECTED ERROR DETECTED BY THE MICF DBID VALIDATION ROUTINE. RETURN CODE IS <retcode>.

EXAMINE THE UNIT DATABASE ID LIST
SPECIFIED OR DEFAULTED FOR THIS REPORT.
REPORT EXECUTION IS TERMINATED.

TYPE: Error

REASON: The MICF routine used to validate unit databases (DBIDs) encountered an unexpected error. In the message text, <retcode> is replaced with the error code returned from MICF.

ACTION: Processing terminates. Validate the unit database ID list used by this report. Make sure it still represents valid CA MICS units in your complex.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 4 6 |  
+-----+
```

TEXT: THE FOLLOWING UNIT DATABASE ID LIST VALUES
CONTAIN NONE OF THE FILES USED BY THIS REPORT:
<dbidlst>.
THE SPECIFIED VALUES WILL BE IGNORED.
REPORT EXECUTION CONTINUES.

TYPE: Warning

REASON: One or more of the values in the unit database ID list specified or defaulted for this report will be ignored, as it contains none of the files used by the report. In the message text, <dbidlst> contains the list of ignored DBIDs.

ACTION: Processing continues. If the ignored DBIDs were supplied by you, remove them from the specification.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 4 7 |  
+-----+
```

TEXT: THE FOLLOWING REQUIRED FILES WERE NOT FOUND IN ANY
OF THE DEFAULT/SPECIFIED UNITS: <filelist>.

CORRECT THE UNIT DATABASE ID LIST TO INCLUDE AT
LEAST ONE UNIT CONTAINING THE SPECIFIED FILES.
REPORT EXECUTION IS TERMINATED.

TYPE: Error

REASON: Files that were required for this report to execute
were not found in the unit database ID list that was
supplied or defaulted. In the message text, <filelist> is
replaced with the names of the file(s) that could not be
located.

ACTION: Processing terminates. Make sure the required
files are installed in at least one unit within
your complex. Modify the DBID specification for
the report.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 4 8 |  
+-----+
```

TEXT: NO CYCLE VALUES HAVE BEEN SPECIFIED FOR THIS
REPORT.
SPECIFY AT LEAST ONE VALID CYCLE VALUE WHEN
SELECTING THIS REPORT FOR EXECUTION.
REPORT EXECUTION IS TERMINATED.

TYPE: Error

REASON: This message is issued by the MICF Production
Reporting option. A valid cycle value was not specified, or
could not be defaulted for this report.

ACTION: Processing terminates. Specify a valid cycle number
when selecting this report for MICF Production Reporting.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 4 9 |  
+-----+
```

TEXT: THIS REPORT ALLOWS ONLY ONE CYCLE VALUE, BUT A

CYCLE RANGE HAS BEEN SPECIFIED OR DEFAULTED.
ONLY THE STARTING CYCLE NUMBER WILL BE USED.
REPORT EXECUTION CONTINUES.

TYPE: Warning

REASON: The report you are executing within the MICF
Production Reporting facility supports data from only one
cycle, but you have specified a range of cycles.

ACTION: The range is ignored, and only the first number
specified will be used. Change the selection specification
to contain only one cycle number.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 0 |  
+-----+
```

TEXT: TEST CASE MEMBER <memname> FOUND IN THE TEST
RESULTS REPORT BUT NOT IN THE TEST CASE LIBRARY.
MAKE SURE THE REPORT FILE YOU SPECIFIED WAS
GENERATED FROM THE TEST CASE LIBRARY YOU
SPECIFIED.
REPORT EXECUTION CONTINUES, BUT THIS TEST CASE
WILL BE DROPPED FROM FURTHER REPORTS.

TYPE: Warning

REASON: A test case member identified in the ACS testing
output report could not be found in the test case library
you specified. In the message text, <memname> is replaced
by the member name that could not be located.

ACTION: The specified member is ignored, and will be
excluded from all future reports. Make sure the test case
library and the ACS output report file you specified to this
report represent the same data. The ACS output file should
have been produced from ACS testing using the test case
library you specified as input.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 1 |  
+-----+
```

TEXT: THE FOLLOWING UNEXPECTED VALUE WAS FOUND IN THE
CYCLE REQUIREMENTS INDICATOR: <ind>.
MAKE SURE NO USER CHANGES HAVE BEEN MADE TO
THIS MICF INQUIRY, THEN CONTACT TECHNICAL
SUPPORT.
REPORT EXECUTION IS TERMINATED.

TYPE: Error

REASON: The report you attempted to execute under the MICF
Production Reporting facility had an invalid value in the
cycle requirements indicator. In the message text, <ind> is
replaced by the value found to be in error.

ACTION: Processing terminates. Make sure the inquiry you
attempted to execute has not been modified.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 2 |  
+-----+
```

TEXT: RESULTS MAY NOT BE MEANINGFUL BECAUSE ONLY ONE
DATE WAS FOUND IN THE INPUT DATA FOR THIS REPORT.
CONSIDER RUNNING THIS REPORT AGAIN WITH A RANGE
OF INPUT CYCLES SPECIFIED.
REPORT EXECUTION CONTINUES.

TYPE: Warning

REASON: The report you have selected to execute is more
meaningful if data representing multiple days is processed,
but only one date was found in the input data you specified.

ACTION: Processing continues, but some of the output
produced may not be meaningful. Specify a range of input
cycles when requesting this report. If a range was
specified, make sure the requested cycles are present and
contain valid data.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 3 |  
+-----+
```

TEXT: ONLY THE FIRST 36,000 DATA SETS WITH THE PREFIX
<dsnpref> WILL BE INCLUDED IN THE OUTPUT.
CONSIDER ALTERNATE INPUT OR FILTERING OPTIONS TO
REDUCE THE AMOUNT OF DATA SELECTED. REPORT
EXECUTION CONTINUES.

TYPE: Warning

REASON: The member name index structure used when building
ACS test cases has a limitation of 36,000 data sets with the
same 3-character prefix, and this limitation has been
exceeded. In the message text, <dsnpref> is the prefix
being processed when the error was detected.

ACTION: Processing continues, but all remaining data sets
with this prefix will be ignored. Change the selection,
filtering, or subsetting values you specified so that fewer
data sets are selected with this prefix. In most cases,
this number of test cases would be unmanageable anyway, and
should be reduced.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 4 |  
+-----+
```

TEXT: A JOBSTREAM TO CREATE AND POPULATE AN ACS TEST
CASE LIBRARY HAS BEEN BUILT IN THE FOLLOWING DATA
SET:
<dsname(member)>
PLEASE REVIEW AND THEN SUBMIT THIS JOBSTREAM TO
BUILD THE ACS TEST CASE LIBRARY.

TYPE: Information

REASON: A jobstream has been successfully created and
placed in your Work Data Set. When this jobstream is
submitted, it will allocate an ACS test case library and

populate it with test cases. In the message text, <dsname> represents your Work Data Set and <member> represents the member created to contain the jobstream.

ACTION: Submit the specified member.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 5 |  
+-----+
```

TEXT: THE FOLLOWING ACS TESTING REPORT FILE WAS
SPECIFIED AS INPUT FOR THIS ANALYSIS:
<dsname>(NONE)

TYPE: Information

REASON: This message documents the name of the ACS testing output report data set that you specified as input to this analysis. The data set should have been created by the ACS testing facility, when testing a test case library you previously built. In the message text, <dsname> is replaced with the name of the report file you supplied.

ACTION: None.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 6 |  
+-----+
```

TEXT: THE FOLLOWING ACS TEST CASE DATA SET WAS SPECIFIED
AS INPUT FOR THIS ANALYSIS:
<dsname>(ZZZINDEX)
THIS DATA SET SHOULD HAVE BEEN CREATED
PREVIOUSLY AND THEN SPECIFIED TO THE ISMF ACS
TESTING FACILITY.

TYPE: Information

REASON: Specifies the ACS test case library that was specified as input to this report. The test case library is needed to cross reference the material in the original test

cases with the information in the testing report. It is important that these two data sets represent the same data. In the message text, <dsname> represents the name of the test case library you specified.

ACTION: None.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 7 |  
+-----+
```

TEXT: THE REQUESTED ANALYSIS REPORT(S) CANNOT BE PRODUCED BECAUSE THE STORAGE CLASS ACS ROUTINE WAS NOT TESTED.
EXECUTE THE ISMF ACS TEST FACILITY AND SPECIFY THAT THE STORAGE CLASS ACS ROUTINE BE TESTED.
PROCESSING CONTINUES, BUT THIS REPORT WILL NOT BE PRODUCED.

TYPE: Warning

REASON: You specified that one or more of the analysis reports should be produced, yet SMS Storage Class was not tested by the ACS testing facility. You have to have specified Storage Class testing in the IBM ACS test facility in order to produce the analysis reports, as they are based on the results of the tests.

ACTION: Processing continues, but the analysis reports will not be produced. Either invoke the ACS test facility again, this time testing the storage class ACS routine, or do not request the analysis reports.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 8 |  
+-----+
```

TEXT: MORE THAN 9 ELEMENTS HAVE BEEN SPECIFIED FOR THE SAME CHART IN PLOT GROUP: <grpid>.
RESPECIFY ELEMENT SELECTION SO THAT A MAXIMUM

OF 9 VALUES APPEAR IN ANY ONE GROUP.
PROCESSING CONTINUES; SOME ELEMENTS WILL BE
DROPPED FROM THE OUTPUT CHARTS.

TYPE: Warning

REASON: SAS permits a maximum of nine elements to be plotted on the same graph, but you have exceeded that limit. You specify elements that should be grouped together by selecting them with the same alphabetic character on the selection panel. In the message text, <grpid> represents the character used to select the elements in this group.

ACTION: Processing continues, but all elements over the stated limit will be ignored. Either specify fewer elements to appear on this graph, or use a different selection character for some elements, so they will appear on a separate graph.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 0 0 5 9 |  
+-----+
```

TEXT: <type> HSM SPACE REPORTING IS IN EFFECT FOR THIS
REPORT EXECUTION.

TYPE: Information

REASON: There are three different types of space values available for data sets under the control of HSM, and this message indicates which of those options will be used. In the message text, <type> contains the value PRIMARY (space allocated on primary volume), ACTUAL (space allocated on the HSM volume), or ESTIMATED (space estimate if moved back to primary). These options can be changed from the StorageMate and MICF Options panel.

ACTION: None.

REFERENCES: CA MICS StorageMate Guide, Section 3.9.

```
+-----+
```

```
| S T G 0 0 0 6 0 |  
+-----+
```

TEXT: MORE THAN 2 MEASUREMENT SAMPLES HAVE BEEN DETECTED
FOR AT LEAST ONE INPUT DATE.
DETAILED REPORTING WILL BE DONE FOR THIS REPORT
EXECUTION.

TYPE: Information

REASON: The report logic has detected more than two
measurement samples for at least one volume on at least one
day. This means that the CA MICS DASD Collector was
probably run multiple times on the same day to collect
storage information at different times throughout the day.

ACTION: The report logic is changed to use detailed
reporting mode, rather than daily mode. Each graph or
report will be used to represent one day with the different
measurement times being used as entries on each page. If
this is not what was desired, use different CA MICS data
without multiple measurements for the same day.

REFERENCES: CA MICS StorageMate Guide, Section 4.3.

```
+-----+  
| S T G 0 2 0 0 0 |  
+-----+
```

TEXT: "STGMSG"

TYPE: Informational

REASON: Multi-use generic message

ACTION: none.

REFERENCES: none.

```
+-----+  
| S T G 0 2 0 1 0 |  
+-----+
```

TEXT: "STGMSG"

TYPE: Informational

REASON: Multi-use generic message; future use

ACTION: none.

REFERENCES: none.

```
+-----+
| S T G 0 2 0 5 0 |
+-----+
```

TEXT: "PHASE": "_PGMTTL" BEGINNING EXECUTION...

TYPE: Informational

REASON: Processing phase _PGMTTL is starting

ACTION: none.

REFERENCES: none.

```
+-----+
| S T G 0 2 0 6 0 |
+-----+
```

TEXT: NEUGENTS MODEL "_FUNCT" PHASE BEGINNING...

TYPE: Informational

REASON: Neugents internal phase _FUNCT starting

ACTION: none.

REFERENCES: None.

```
+-----+
| S T G 0 2 0 7 0 |
+-----+
```

TEXT: "PHASE": "_PGMTTL" ENDING EXECUTION...

TYPE: Informational

REASON: Processing phase _PGMTTL ending.

ACTION: none.

REFERENCES: None.

```
+-----+  
| S T G 0 2 0 7 5 |  
+-----+
```

TEXT: "PHASE": "_PGMTTL" ABNORMALLY TERMINATING

TYPE: Error.

REASON: Processing _PGMTTL abnormally terminated.

ACTION: Neugents processing job will end.

REFERENCES: None.

```
+-----+  
| S T G 0 2 0 8 0 |  
+-----+
```

TEXT: "V1" FILE CONTAINS "N" RECORDS SELECTED FOR
PROCESSING.

TYPE: Informational

REASON: N records have been selected for processing
from file V1.

ACTION: none.

REFERENCES: None.

```
+-----+  
| S T G 0 2 0 8 5 |  
+-----+
```

TEXT: DATA FROM FILE "V1" WILL BE EXCLUDED, YOUR
RESULTS MAY BE NEGATIVELY AFFECTED.

TYPE: Warning.

REASON: Some observations from file VI were excluded,
either due data value errors or unmatched

observations among files. Modeling results may be impacted and should be examined in light of this message.

ACTION: none.

REFERENCES: None.

```
+-----+  
| S T G 0 2 0 9 0 |  
+-----+
```

TEXT: NO FORECAST FILE ENTRY FOUND FOR "V1" ENTRY:

TYPE: Warning.

REASON: An observation in the Primary STG file was not matched in the alternate input file.

ACTION: The original observation is not used in further processing. See message STG02095 for additional details.

REFERENCES: None.

```
+-----+  
| S T G 0 2 0 9 5 |  
+-----+
```

TEXT: %FNAME %FDATA

TYPE: Warning.

REASON: One of the key values of the unmatched observation is displayed. FNAME is the name of the key variable; FDATA is the data value of the key variable.

ACTION: None.

REFERENCES: None.

```
+-----+  
| S T G 0 2 1 0 0 |  
+-----+
```

TEXT: REPORT: "V1" BEING CURRENTLY GENERATED ON
DDNAME //"V1"

TYPE: Informational

REASON: Generation of report V1 has started.

ACTION: The report will be written out on DDNAME V1.
This message will only be issued if the MICF
environment is not being used. Message STG02105
will normally be issued instead.

REFERENCES: None.

+-----+
| S T G 0 2 1 0 5 |
+-----+

TEXT: REPORT: "V1" BEING CURRENTLY GENERATED.

TYPE: Informational

REASON: Generation of report V1 has started.

ACTION: Reports generated within the MICF environment
are normally written out to the SASLIST DD.

REFERENCES: None.

+-----+
| S T G 0 2 1 1 0 |
+-----+

TEXT: REPORT: "V1" BEING CURRENTLY GENERATED VIA CSV
FILE.

TYPE: Informational

REASON: Generation of report V1 has started and will be
created using a CSV file.

ACTION: This option is used to create MQR compatible
CSV files.

REFERENCES: None.

```
+-----+  
| S T G 0 2 1 2 0 |  
+-----+
```

TEXT: GRAPH: "V1" BEING CURRENTLY GENERATED ON DDNAME
// "V1"

TYPE: Informational

REASON: Generation of graph V1 has started.

ACTION: The graph will be written out on DDNAME V1.
This message will only be issued if the MICF
environment is not being used. Message STG02125
will normally be issued instead.

REFERENCES: None.

```
+-----+  
| S T G 0 2 1 2 5 |  
+-----+
```

TEXT: GRAPH: "V1" BEING CURRENTLY GENERATED.

TYPE: Informational

REASON: Generation of graph V1 has started.

ACTION: Graphs generated within the MICF environment
are normally written out to the SASLIST DD.

REFERENCES: None.

```
+-----+  
| S T G 0 2 1 3 0 |  
+-----+
```

TEXT: GRAPH: "V1" BEING CURRENTLY GENERATED VIA CSV
FILE.

TYPE: Informational

REASON: Generation of graph V1 has started and will be
created using a CSV file.

ACTION: This option is used to create MQR compatible
CSV files.

REFERENCES: None.

```
+-----+  
| S T G 0 2 1 5 0 |  
+-----+
```

TEXT: INVALID DATA VALUES DETECTED IN "V1" FILE,
"_PGMTTL" CANNOT PROCESS SUCCESSFULLY, WHEN
DATA VALUES ARE MISSING OR INVALID. SEE
BADTRAIN REPORT FOR DETAILS.

TYPE: Warning

REASON: Missing or zero values were detected in data
elements in one or more observations.

ACTION: The observations are not used in further
processing. Report "BADTRAIN" will contain a
list of observations that were dropped.

REFERENCES: None.

```
+-----+  
| S T G 0 2 1 6 0 |  
+-----+
```

TEXT: MODEL SUMMARY STATISTICS HAVE BEEN CREATED ON
DDNAME //MODLSTAT

TYPE: Informational

REASON: The Model Summary Statistics report is being
created.

ACTION: The report will be written out on DDNAME
MODLSTAT. This message will only be issued if
the MICF environment is not being used. Message
STG02165 will normally be issued instead.

REFERENCES: None.

```
+-----+  
| S T G 0 2 1 6 5 |  
+-----+
```

TEXT: MODEL SUMMARY STATISTICS HAVE BEEN CREATED.

TYPE: Informational

REASON: The Model Summary Statistics report is being created.

ACTION: Reports generated within the MICF environment are normally written out to the SASLIST DD.

REFERENCES: None.

```
+-----+  
| S T G 0 2 1 7 0 |  
+-----+
```

TEXT: "_FORTYP" FORECASTING OPTION SELECTED -, "V1" "V2" OF PREDICTED DATA WILL BE GENERATED. BEGINNING WITH "X_SDATE".

TYPE: Informational.

REASON: For forecasting option _FORTYP, V1 months/weeks/days of data were selected for processing, beginning with X_DATE. This message defined the time frame of the modeling study.

ACTION: None.

REFERENCES: None.

```
+-----+  
| S T G 0 2 3 0 0 |  
+-----+
```

TEXT: MACRO "PHASE" BEGINNING EXECUTION...

TYPE: Informational

REASON: Macro PHASE is starting execution. This macro is used to provide linkage to Neugents technology.

ACTION: None.

REFERENCES: None.

```
+-----+  
| S T G 0 2 3 1 0 |  
+-----+
```

TEXT: MACRO "PHASE" FAILED WITH RETURN CODE "RC"

TYPE: Error.

REASON: A non-zero return code value was detected upon
 return from Neugents.

ACTION: Neugents processing will be terminated.

REFERENCES: None.

```
+-----+  
| S T G 0 2 3 2 0 |  
+-----+
```

TEXT: MACRO GLOBALS MISSING REQUIRED OPERAND
 "MODNAME", EXECUTION TERMINATING.

TYPE: Error.

REASON: An internal error has occurred. This message
 should never occur except during product
 development.

ACTION: Processing terminates.

REFERENCES: None.

```
+-----+  
| S T G 0 2 3 2 5 |  
+-----+
```

TEXT: MACRO GLOBALS MISSING REQUIRED OPERAND
 "MODTITLE", EXECUTION TERMINATING.

TYPE: Error

REASON: An internal error has occurred. This message
 should never occur except during product
 development.

ACTION: Processing terminates.

REFERENCES: None.

```
+-----+  
| S T G 0 2 3 9 0 |  
+-----+
```

TEXT: MACRO "PHASE" TERMINATING, ABORT REQUESTED.

TYPE: Error.

REASON: A terminal error condition was detected in
macro PHASE. This message is preceded by
either message STG02310 or STG02590.

ACTION: Neugents processing terminates.

REFERENCES: None.

```
+-----+  
| S T G 0 2 3 9 5 |  
+-----+
```

TEXT: MACRO "PHASE" ENDED WITH RETURN CODE "RC"

TYPE: Informational

REASON: The macro PHASE ended execution, with return
code RC.

ACTION: None.

REFERENCES: None.

```
+-----+  
| S T G 0 2 4 1 0 |  
+-----+
```

TEXT: MODEL TRAINING PHASE COMPLETED SUCCESSFULLY.
TRAINING STATISTICS ARE SUMMARIZED ON DDNAME
//DELTASTS, AND ARE INCLUDED BELOW: NUMBER OF
CASES USED FOR TRAINING: "N",
AVERAGE MODEL ERROR MEASURED: "V1",
AVERAGE MODEL RMS ERROR MEASURED: "V2",
(GOODNESS OF FIT) FOR MODEL: "V3"

TYPE: Informational

REASON: Neugents Value Prediction models first execute a training phase, and produce several statistics used to evaluate the model. These statistics are displayed in the message.

ACTION: The client should review the statistics before deciding on the validity of the model.

REFERENCES: None.

```
+-----+  
| S T G 0 2 4 2 0 |  
+-----+
```

TEXT: MODEL FORECASTING PHASE COMPLETED SUCCESSFULLY.
FORECASTING STATISTICS ARE SUMMARIZED ON DDNAME
//CONSFCS, AND ARE INCLUDED BELOW: NUMBER OF
CASES USED FOR FORECASTING: "N",
(GOODNESS OF FIT) FOR MODEL: "V1"

TYPE: Informational

REASON: Neugents Value Prediction models secondly execute a consulting phase, and produce additional statistics used to evaluate the model. These statistics are displayed in the message.

ACTION: The client should review the statistics before deciding on the validity of the model.

REFERENCES: None.

```
+-----+  
| S T G 0 2 4 3 0 |  
+-----+
```

TEXT: LINEAR REGRESSION PHASE COMPLETED SUCCESSFULLY.
MODEL STATISTICS ARE SUMMARIZED ON DDNAME
//MODLSTAT, AND ARE INCLUDED BELOW: NUMBER OF
CASES USED FOR MODELING: "V1",
SLOPE OF REGRESSION LINE: "V2",

INTERCEPT VALUE FOR IND. ELEMENT: "V3",
(GOODNESS OF FIT) FOR MODEL: "V4"

TYPE: Informational

REASON: Workload Forecasting will execute a simple linear regression operation when a single file is specified and only a forecasting element is provided. Regression models produce various statistics used to evaluate the model. These statistics are displayed in the message.

ACTION: The client should review the statistics before deciding on the validity of the model. The report will be written out on DDNAME MODLSTAT. This message will only be issued if the MICF environment is not being used. Message STG02435 will normally be issued instead.

REFERENCES: None.

```
+-----+  
| S T G 0 2 4 3 5 |  
+-----+
```

TEXT: LINEAR REGRESSION PHASE COMPLETED SUCCESSFULLY.
MODEL STATISTICS ARE SUMMARIZED ON DDNAME
//SASLIST AND ARE INCLUDED BELOW: NUMBER OF
CASES USED FOR MODELING: "V1",
SLOPE OF REGRESSION LINE: "V2",
INTERCEPT VALUE FOR IND. ELEMENT: "V3",
(GOODNESS OF FIT) FOR MODEL: "V4"

TYPE: Informational

REASON: Workload Forecasting will execute a simple linear regression operation when a single file is specified and only a forecasting element is provided. Regression models produce various statistics used to evaluate the model. These statistics are displayed in the message.

ACTION: The client should review the statistics before deciding on the validity of the model. Reports

generated within the MICF environment are normally written out to the SASLIST DD.

REFERENCES: None.

```
+-----+  
| S T G 0 2 5 0 0 |  
+-----+
```

TEXT: OPEN ISSUED FOR "STD" FILE. RETURN CODE IS:
"RC"

TYPE: Informational

REASON: An open was issued for file with DDNAME "STD". Such files are used by Neugents processing for passing data between Neugents and the Capacity Planning applications, or are required by USS for internal use. The return code from open processing is displayed as "RC".

ACTION: None.

REFERENCES: None.

```
+-----+  
| S T G 0 2 5 1 0 |  
+-----+
```

TEXT: OPEN FAILED FOR "STD" FILE. SYSTEM ERROR
MESSAGE IS "SYSMSG" EXECUTION TERMINATED.

TYPE: Error

REASON: An open was issued for file with DDNAME "STD". Such files are used by Neugents processing for passing data between Neugents and the Capacity Planning applications, or are required by USS for internal use. The return code from open processing is displayed as "RC". Generally, a non-zero return code value will cause processing to be terminated.

ACTION: Review the specific system message for action.

REFERENCES: None.

```
+-----+  
| S T G 0 2 5 9 0 |  
+-----+
```

TEXT: ERROR TEXT FROM NEUGENTS PROCESSING "ERRMSG"
EXECUTION TERMINATED.

TYPE: Error

REASON: Any error text returned by USS in the STDERR
dataset will be displayed herein.

ACTION: Review specific error message for action.

REFERENCES: None.

Appendix B: INSTALLATION

This appendix describes the installation process for CA MICS StorageMate and is intended for use by the CA MICS System Administrator.

This appendix contains a checklist describing the installation steps that are unique to installing the CA MICS StorageMate product.

Installation Checklist

- ___ 1. If you are not coming to this checklist from one of the checklists in Chapter 3 of the PIOM, or from the StorageMate Product Installation Checklist, STG0000, exit this procedure now.

+-----+
| Reference(s): Chapter 3, CA MICS PIOM (new |
| CA MICS) STG0000 (optional product |
| STG0000 (optional product only) |
+-----+

- ___ 2. Ensure that you are authorized to allocate data sets.

+-----+
| Reference(s): Section 2.1 |
+-----+

- ___ 3. Ensure that the required files and their timespans are available from the VCA, SMF, and RMF products. If you plan to use the optional ASTEX and HSM files, make sure they are also available.

+-----+
| Reference(s): Section 2.2 |
+-----+

- ___ 4. Define CA MICS StorageMate processing option macros.

+-----+
| Reference(s): Section 3.2 |
+-----+

___ 5. If installing the product from the PIOM, exit this procedure and return to the next step in the PIOM.

___ 6. Identify the CA MICS unit databases that contain the files needed by CA MICS StorageMate.

```
+-----+
| Reference(s):  Section 2.2          |
+-----+
```

```
* * * * *
*
* The remainder of this procedure, steps 7 through 11,
* must be completed by each person who intends to use
* this product. They do not have to be completed by the
* person who is installing the product unless that person
* intends to use the product as well.
*
* * * * *
```

___ 7. Execute a sample inquiry.

```
+-----+
| Reference(s):  Section 3.3          |
+-----+
```

___ 8. If desired, modify your MICF options.

```
+-----+
| Reference(s):  Sections 3.6, 3.7, and 3.8          |
+-----+
```

___ 9. Create a sample Volume Group Table.

```
+-----+
| Reference(s):  Sections 3.4, 3.4.1, 3.4.2          |
+-----+
```

'__ 10. Create a sample Data Set Group Table.

```
+-----+
| Reference(s):  Sections 3.5, 3.5.1, 3.5.2          |
+-----+
```

___ 11. Become familiar with data filtering.

```
+-----+
| Reference(s):  Section 3.9      |
+-----+
```

Installation is now complete. Please review the remaining sections of this guide and explore the features of CA MICS StorageMate.

Appendix C: TAPE INTERFACE

While information about your tape library does not currently reside in CA MICS, CA MICS StorageMate provides an external interface that allows you to place tape-related information in a non-CA MICS SAS data set. This section documents the steps you must take to create this external data set before attempting to use tape data in any of the CA MICS StorageMate reports. You do not need to follow these steps unless you plan to process TMC or TLMS data as an input source to CA MICS StorageMate.

Five sample jobs are distributed with this product as an example of how external tape library information can be placed into a SAS data set and then read as input by CA MICS StorageMate. These sample jobs are contained in the following members:

- o sharedprefix.MICS.CNTL(STGEXTP1) -- for CA-1 4.9
- o sharedprefix.MICS.CNTL(STGEXTP3) -- for CA-1 5.0
- o sharedprefix.MICS.CNTL(STGEXTP4) -- for CA-1 5.2 and higher
- o sharedprefix.MICS.CNTL(STGEXTP2) -- for TLMS 5.4
- o sharedprefix.MICS.CNTL(STGEXTP5) -- for TLMS 5.5 and higher

The sample jobs are designed to use data from your tape management software system as input into the CA MICS StorageMate product. The distributed samples are designed to use data from two of the more popular tape management library systems. Other tape library software can be used if you can adapt the samples to process the tape data correctly. The documentation included with the sample jobs describes the name of the SAS data set you must create, and the elements it must contain. Any routine that creates a data set in this described format should be able to work with the CA MICS StorageMate product.

The first step in the STGEXTP1, STGEXTP3 and STGEXTP4 sample jobs extracts tape information from the tape master catalog using the TMSDATA utility supplied with TMS/CA-1 and writes it to a flat file. To convert the data to SAS format, a second step is executed that converts the flat file to a SAS database file with one observation present for each active tape.

The STGEXTP2 and STGEXTP5 sample jobs read data directly from the TLMS VMF file and convert it into a SAS data set, also with each tape represented by one observation.

Macro %STGTAPEX within each sample job may need to be modified to set the following element correctly:

TAPACTx - Similar to other CA MICS files, up to nine account fields may be defined and populated. By default, all sample routines define three accounting elements (TAPACT1-TAPACT3) and populate them based on the accounting information in the raw TMS and TLMS accounting fields, as follows:

TAPACT1 - Positions 1-5 of the accounting area
TAPACT2 - Positions 6-10 of the accounting area
TAPACT3 - Positions 11-15 of the accounting area

For more detailed information, refer to the provided samples.

Appendix D: CASE STUDIES

CASE STUDY #1

The ABC Soup Company has used Generation Data Groups (GDGs) on tape for many years, but has only recently allowed GDGs to be created on DASD. The Storage Administrator has been recently assigned to analyze the current DASD-resident GDGs, to write some basic standards for these types of data sets, to identify any current problems, and to notify the users responsible for those problems.

The reports produced by the GDG Management and Retention facility of CA MICS StorageMate can help to perform this study. The Storage Administrator invokes this facility by using the following panel selection sequence:

- o Display the CA MICS StorageMate primary option panel and select the 'Storage Administration' option
- o Select the 'Data Set Standards' option
- o Select the 'GDG Management and Retention' facility

At this point the execution options panel for this facility is displayed. The Storage Administrator completes the options on the panel so that it appears as shown in Figure D-1.

```

----- GDG Management and Retention -----
Command ==>

CA MICS Input Definition:
  Unit DBID(s) ==> V          Cycle ==> 01 Sources ==> D

Report Options:
  Produce prefix report      ==> N      (Y/N)
  Produce GDG base report   ==> Y      (Y/N)
  Produce detail report     ==> N      (Y/N)
  Simulate generations kept ==> ____   (Numeric, 1-999)
  For prefix/GDG base reports:
    Sort reports by         ==> DSNAME (DSNAME/SAVINGS)

DOWN - Filter  ENTER - Validate END - Run
-----

```

Figure D-1. Execution Options Panel

The input data for the report comes from the 01 data cycle in the 'V' Unit of the CA MICS database. The 'D' in the Sources area indicates that only DASD data should be used. The ISPF DOWN key could be used to display the data filtering panel, but in this case no special data filtering is done. Only one of three possible reports (the GDG base report), is produced and the data on that report is sorted based on the GDG data set name.

The Storage Administrator completes the panel as shown above, and produces the report. A copy of the report output appears in Figure D-2.

GDG MANAGEMENT AND RETENTION (STGEFE) CA MICS StorageMate - GDG LEVEL REPORT												
----- System Identifier=SYSA -----												
Base GDG Name	# of GENS	Min GEN #	Max GEN #	Dupe GENS	# of Gaps	# of Vols	Min MB Alloc	Max MB Alloc	Avg MB Alloc	Total MB Alloc	+-Simulated Data Sets Moved	Space Savings-+ MB Saved
BACKUP.DAILY.LIST	5	4	8	0	0	1	0.0	0.0	0.0	0	0	0.0
BACKUP.DATASETS.LIST	15	994	1008	0	0	1	0.1	0.1	0.1	2	0	0.0
BACKUP.DELETED.MVS340.LIST	20	181	200	0	0	1	0.0	0.2	0.2	3	0	0.0
BUDGT.B003.OUTPUT	7	16	22	0	0	4	0.1	0.2	0.2	1	0	0.0
BUDGT.TEST.GDG	4	1	4	0	0	3	0.1	2.2	0.6	2	0	0.0
COS.PROD.DB5.BACKUP.SDLDS	2	516	516	1	0	2	0.1	0.1	0.1	0	0	0.0
DEVPROD.REPORTS	2	1	2	0	0	2	0.1	0.2	0.1	0	0	0.0
DEVTEST.SAS.DAY.DATA	1	18	18	0	0	1	5.9	5.9	5.9	6	0	0.0
DEVTEST.SAS.MONTH.DATA	1	1	1	0	0	1	2.5	2.5	2.5	3	0	0.0
DEVTEST.SAS.WEEK.DATA	1	1	1	0	0	1	2.5	2.5	2.5	3	0	0.0
DUMMY10.TEST.GDG1	1	1	1	0	0	1	0.0	0.0	0.0	0	0	0.0
DZZ110.SMFDAILY.DATA	1	1	1	0	0	1	6.4	6.4	6.4	6	0	0.0
DZZ199.SMFDAILY.DATA	21	85	104	1	0	2	0.7	32.3	7.2	151	0	0.0
XXKUP.DATASETS.LIST	1	918	918	0	0	1	0.5	0.5	0.5	0	0	0.0

SYSID	82			2	0					178	0	0.0

Figure D-2. Sample Report Output

This report shows that 82 generation data sets currently reside on DASD, belonging to slightly more than a dozen GDGs. Several groups have only one generation, while the largest group has 21. The report also identifies several potential problems that should be investigated:

- o The BACKUP.DELETED.MVS340.LIST GDG consists of 20 different generations, but the '# OF VOLS' column indicates that all of these data sets reside on one volume. This could cause a problem if that volume were lost. The JCL for that GDG should be changed to spread the generations across several volumes for better disaster recovery. All three of the GDGs with the BACKUP prefix have this same problem.
- o The COS.PROD.DB5.BACKUP.SDLDS generation contains two data sets, but the minimum and maximum generation numbers found were both 516, and the 'DUPE GENS' column contains a value of 1. This indicates that two data sets exist with the same generation number, 516. Because duplicate data sets cannot both be cataloged, one of these is probably in error and could be deleted. The Storage Administrator could run this report again, this time requesting the detail option and selecting only the COS prefix. This would give more information about both data sets, including the volumes on which they reside.

- o The DZZ199.SMFDAILY.DATA GDG also has one duplicate generation, and thus has the same problem described previously. It could also be investigated and corrected by running the report again with the detail option and requesting only the DZZ199 prefix.
- o Several of the GDGs with DEVTEST and DZZ110 prefixes have only one generation, and that generation number is 1. It is possible these data sets were established but then never used, or were created as a test but then never deleted.
- o The DZZ199.SMFDAILY.DATA group contains 21 generations with an average size of 7.2 megabytes per generation, and at least one generation that is using 32.3 megabytes. Perhaps space could be saved by reducing the number of generations to be kept for this group.
- o The entry XXCKUP.DATASETS.LIST is probably not a real GDG at all, but a renamed version of generation 918 of the BACKUP.DATASETS.LIST group. Because the oldest generation in that group is generation 994, this data set is at least 76 generations old, and can probably be deleted.

Using this report, the Storage Administrator has detected several problems that should be corrected, and has a feel for the standards that should be put in place for DASD-resident Generation Data Groups. One of these standards could be the specification that DASD-resident GDGs should have a maximum of only seven kept generations. A glance at the report shows that a new standard such as this would require changes to only three existing GDGs.

As part of the justification for a new standard, the Storage Administrator wants to calculate the space savings if the standard were implemented. This can be done by executing the GDG Management and Retention facility again with slightly different options. The facility is accessed again in the same manner, but this time the execution options panel will be completed as in Figure D-3.

```

----- GDG Management and Retention -----
Command ==>

CA MICS Input Definition:
  Unit DBID(s) ==> V          Cycle ==> 01 Sources ==> D

Report Options:
  Produce prefix report      ==> Y      (Y/N)
  Produce GDG base report    ==> N      (Y/N)
  Produce detail report      ==> N      (Y/N)
  Simulate generations kept  ==> 7__    (Numeric, 1-999)
  For prefix/GDG base reports:
    Sort reports by          ==> SAVINGS (DSNAME/SAVINGS)

DOWN - Filter  ENTER - Validate END - Run
-----

```

Figure D-3. Execution Options Panel

For this execution, the Storage Administrator has requested that the Prefix report be produced, rather than the GDG base report. This will provide the same information, but will summarize at the data set prefix level, rather than at the GDG level. A simulation value of seven was also specified, so that the report will simulate the savings of keeping a maximum of seven generations for each GDG. Finally, the report should be sorted so that prefixes showing the greatest amount of simulated savings will appear first.

The report produced by the previous panel is shown in Figure D-4.

GDG MANAGEMENT AND RETENTION (STGEFE) CA MICS StorageMate - PREFIX REPORT														
----- System Identifier=SYSA -----														
Prefix	# Of GDGs	Min # Gen	Max # Gen	Avg # Gen	GDGs with Dups	GDGs with Gaps	Min # Vol	Max # Vol	Avg # Vol	Max MB Alloc	Avg MB Alloc	Total MB Alloc	+--- Data Sets Moved	Simulated Space Savings MB Saved
DZZ199	1	21	21	21	1	0	2	2	2	151.2	7.2	151.2	14	81.9
BACKUP	3	5	20	13	0	0	1	1	1	3.1	0.1	5.3	21	3.0
BUDGT	2	4	7	6	0	0	3	4	4	2.4	0.3	3.5	0	0.0
COS	1	2	2	2	1	0	2	2	2	0.2	0.1	0.2	0	0.0
DEVPROD	1	2	2	2	0	0	2	2	2	0.2	0.1	0.2	0	0.0
DEVTEST	3	1	1	1	0	0	1	1	1	5.9	3.7	11.0	0	0.0
DUMMY10	1	1	1	1	0	0	1	1	1	0.0	0.0	0.0	0	0.0
DZZ110	1	1	1	1	0	0	1	1	1	6.4	6.4	6.4	0	0.0
XXCKUP	1	1	1	1	0	0	1	1	1	0.5	0.5	0.5	0	0.0
-----	14				2	0						178.4	35	84.9

Figure D-4. Sample Report Output

This report shows that only two groups would be affected by the proposed standard of seven kept generations -- prefix DZZ199 and prefix BACKUP. Implementation of this standard would cause 35 data sets to be removed, resulting in space savings of 84.9 megabytes. The Storage Administrator can now use these values to justify the new standard, and can contact the owners of the two identified prefixes to begin implementation.

CASE STUDY #2

The Maple Manufacturing Company implemented DFSMS about six months ago. At that time, it defined separate Storage Groups for TSO data sets and Batch data sets, as those are their two most important applications. The Storage Administrator wants to monitor these groups on a regular basis, to ensure that they deliver the desired performance and availability. This monitoring function can easily be done using the Volume Group Configuration facility of CA MICS StorageMate. This facility can be invoked by using the following panel selection sequence:

- o Display the CA MICS StorageMate primary option panel and select the 'Storage Administration' option
- o Select the 'Volume Grouping' option

o Select the 'Volume Group Configuration' facility

At this point the execution options panel for this facility is displayed. The Storage Administrator completes the options on the panel so it appears as shown in Figure D-5.

```

----- Volume Group Configuration -----
Command ==>

CA MICS Input Definition:
  Unit DBID(s) ==> P          Cycles ==> 01 - 01

Report Options:
  Produce reports   ==> DETAIL  (SUMMARY/DETAIL/BOTH)
  Volume Group Table ==> _____ (Input member name)

DOWN - Filter  ENTER - Validate END - Run
-----

=====

----- StorageMate Data Filtering -----
Command ==>

Element Oper Value(s)
DSNAME    EQ  _____
GROUP     EQ  TSO  BATCH
VOLSER    EQ  _____

UP - Report    ENTER - Validate
-----

```

Figure D-5. Execution Options Panel

The input data for the report comes from the 01 data cycle in the 'P' Unit of the CA MICS database. Data is filtered so that only the Storage Groups named TSO and BATCH will appear on the report. The Detailed version of the output report has also been requested. The example shown above is actually a portion of two different panels. The DOWN key is used to display the filtering panel, and then the UP key is used to return to the main report panel.

The Storage Administrator completes the panels as shown above, and produces the report. A copy of the report output appears in Figure D-6.

VOLUME GROUP CONFIGURATION (STGEHA) CA MICS StorageMate - DETAIL														
----- System Identifier=SYSB / Volume Group=BATCH -----														
Device Address	C.U. Path	Volume Serial Number	Device Type	Frag. Index	+---- Total	Data Sets Non- VSAM	----- VSAM	+-- Total	Tracks Allocated Non- VSAM	--- VSAM	+-- Free Tracks	--- Per- cent	+-- Average I/O Rate	-- Response Time
140	0001	STG140	3380	0.314	60	59	1	8069	8066	3	5061	38%	4.43	24.4
141	0001	STG141	3380	0.052	117	105	12	11148	9853	1295	2096	16%	4.47	24.8
142	0001	STG142	3380	0.099	6	1	5	9318	15	9303	3896	29%	2.82	32.0
143	0001	STG143	3380	0.948	351	348	3	12876	12710	166	338	3%	0.73	22.4
144	0001	STG144	3380	0.375	134	134	0	11541	10142	1399	1703	13%	0.00	26.6
145	0001	STG145	3380	0.312	2	2	0	13245	13245	0	24	0%	4.01	24.0
146	0001	STG146	3380	0.463	352	345	7	12963	12823	140	251	2%	0.75	20.5
147	0001	STG147	3380	0.805	248	245	3	12679	12358	321	535	4%	0.70	21.3
148	0001	STG148	3380	0.338	139	130	9	12933	12718	215	311	2%	0.23	22.1
149	0001	STG149	3380	0.399	64	63	1	12286	12283	3	958	7%	1.19	21.3
14A	0001	STG14A	3380	0.833	229	218	11	13121	12501	620	30	0%	0.74	19.9
14B	0001	STG14B	3380	0.793	202	201	1	12850	12847	3	364	3%	0.51	22.9
14C	0001	STG14C	3380	0.321	207	200	7	13093	4748	8345	121	1%	1.34	23.2
14D	0001	STG14D	3380	0.689	198	197	1	12750	12735	15	464	3%	0.59	20.5
14E	0001	STG14E	3380	0.927	236	234	2	13160	13123	37	54	0%	0.48	18.2
14F	0001	STG14F	3380	0.803	345	339	6	12456	12327	129	758	6%	0.37	20.6
-----				-----										
016 Volumes				0.529	2890	2821	69	194488	172494	21994	16964	8%	23.38	24.3
----- System Identifier=SYSB / Volume Group=TS0 -----														
Device Address	C.U. Path	Volume Serial Number	Device Type	Frag. Index	+---- Total	Data Sets Non- VSAM	----- VSAM	+-- Total	Tracks Allocated Non- VSAM	--- VSAM	+-- Free Tracks	--- Per- cent	+-- Average I/O Rate	-- Response Time
240	0002	STG240	3380	0.337	136	131	5	12143	10744	1399	1101	8%	1.60	24.7
241	0002	STG241	3380	0.652	159	149	10	12521	12034	487	693	5%	0.71	15.7
242	0002	STG242	3380	0.131	6	1	5	7398	15	7383	5816	44%	2.52	41.1
243	0002	STG243	3380	0.461	224	207	17	12965	12162	803	249	2%	0.86	22.2
244	0002	STG244	3380	0.502	576	541	35	13111	12147	964	103	1%	0.20	33.9
245	0002	STG245	3380	0.382	257	252	5	13055	12658	397	159	1%	0.57	23.9
246	0002	STG246	3380	0.665	213	212	1	12291	12276	15	923	7%	0.26	32.1
247	0002	STG247	3380	0.259	217	212	5	12678	12511	167	536	4%	0.44	21.3
248	0002	STG248	3380	0.521	39	38	1	12745	12742	3	469	4%	0.16	30.7
249	0002	STG249	3380	0.443	227	226	1	12908	12893	15	306	2%	0.29	21.4
24A	0002	STG24A	3380	0.324	120	117	3	12943	12895	48	301	2%	0.50	27.9
24B	0002	STG24B	3380	0.590	234	221	13	13051	12250	801	154	1%	0.77	25.6
24C	0002	STG24C	3380	0.162	216	185	31	12844	11207	1637	370	3%	1.51	23.5
24D	0002	STG24D	3380	0.674	226	225	1	12283	12268	15	931	7%	0.38	29.7
24E	0002	STG24E	3380	0.346	146	143	3	11713	11455	258	1531	12%	0.00	31.9
24F	0002	STG24F	3380	0.044	9	1	8	6104	15	6089	7140	54%	7.65	15.5
-----				-----										
016 Volumes				0.406	3005	2861	144	190753	170272	20481	20782	10%	18.42	22.9

Figure D-6. Sample Report Output

The results from the report indicate that the two groups are similar, in that each contains 16 volumes and has similar fragmentation, usage, free space, and activity values.

Based on projections of future activity, the Storage Administrator expects a greater demand for space from Batch users over the next few months, and a corresponding decrease in the space needed by TSO users. Because of this, the Storage Administrator would like to shift at least two volumes from the TSO group to the batch group. The report shows that volumes STG242 and STG24F would be good candidates, as they contain a total of only 15 TSO data sets, and have plenty of free space.

The Storage Administrator knows that the Volume Group Configuration facility can also be used to simulate alternate groupings and that performing a simulation could be valuable before making the actual volume group changes. Before running the report again, the Administrator would first create a Volume Group Table that would contain SMS volumes (See Chapter 3.4), and then would modify it so that volumes STG242 and STG24F would be assigned to group BATCH, rather than to group TSO. The Administrator would then invoke the Volume Group Configuration facility again, using the options shown in Figure D-7.

```

----- Volume Group Configuration -----
Command ==>

CA MICS Input Definition:
  Unit DBID(s) ==> P          Cycles ==> 01 - 01

Report Options:
  Produce reports    ==> SUMMARY (SUMMARY/DETAIL/BOTH)
  Volume Group Table ==> VGTEST01 (Input member name)

DOWN - Filter  ENTER - Validate END - Run
-----

=====

----- StorageMate Data Filtering -----
Command ==>

Element Oper Value(s)
DSNAME    EQ  -----
GROUP     EQ  TSO  BATCH
VOLSER    EQ  -----

UP - Report    ENTER - Validate
-----

```

Figure D-7. Volume Group Configuration Facility

For this execution, a Volume Group Table named VGTEST01 has been specified. When this table was created, SMS volumes were included in it, and an option was activated that would cause it to be searched for all volumes. This will cause the facility to use the provided table for volume assignments, and ignore the SMS Storage Group assignment. Also, a Summary report was requested this time, instead of the previously shown Detail report.

The report produced by these options is shown in Figure D-8.

VOLUME GROUP CONFIGURATION (STGEHA) CA MICS StorageMate - SUMMARY												
----- System Identifier=SYSB -----												
Volume Group Name	Number of Volumes	Frag. Index	+----- Total	Data Sets Non- VSAM	-----+ VSAM	+----- Total	Tracks Allocated Non- VSAM	-----+ VSAM	+----- Free Tracks	-----+ Per- cent	+--- Average ---+ I/O Rate	---+ Resp. Time
BATCH	18	0.480	2905	2823	82	207990	172524	35466	29920	13%	33.55	23.6
TS0	14	0.451	2990	2859	131	177251	170242	7009	7826	4%	8.24	24.2
-----	32	0.468	5895	5682	213	385241	342766	42475	37746	9%	41.79	23.7

Figure D-8. Sample Report Output

As expected, the number of volumes in each group is no longer 16; two volumes have shifted from the TS0 group to the BATCH group. Note that the percentage of the BATCH group that is unallocated jumped from 8% to 13%, while the free percentage of the TS0 group dropped from 10% to 4%. One drawback of making this change, however, would be to badly skew the average I/O rate to each group, as the BATCH group would receive 33.55 I/Os per second, while the TS0 group receives only 8.24. By making slight changes to the Volume Group Table, the Storage Administrator could run the report several times, using different combinations of volumes, until a better combination was found.

While the I/O rate after this change would be somewhat skewed, the Storage Administrator is satisfied with the other results, and believes that an I/O rate of 33.55 would not have too adverse an effect on the Batch workload. The Administrator decides to implement this change and to run the Volume Group Configuration report on a regular basis to ensure the expected results are achieved.

CASE STUDY #3

XYZ Phone Sales is a small company that just purchased its first cache controller. The main purpose for the new controller is to provide optimum service to their online order entry application. All of the critical data sets related to this application have been moved to volumes under control of this new controller.

The most critical time for this application is Friday morning from 9 a.m. until noon. That is the time their area sales force calls in their orders for the week, and fast response times are a necessity. This critical time window has been defined to their CA MICS system as zone 1.

Now that the new controller is fully implemented, the storage administrator wishes to monitor its activity. This can be done using the ASTEX / Top Resource Consumers report within CA MICS StorageMate. The administrator invokes this facility with the following procedure:

- o Display the CA MICS StorageMate primary option panel and select the 'Performance Management' option
- o Select the 'ASTEX / Top Resource Consumers' option

At this point the execution options panel for this facility is displayed. The Storage Administrator completes the options on the panel so it appears as shown in Figure D-9.

```

----- ASTEX / Top Resource Consumers -----
Command ==>

CA MICS Input Definition:
  Unit DBID(s) ==> A          Cycles ==> 01 - 03

Report Options:
  Produce reports          ==> CHART    (CHART/LIST/BOTH or DETAIL)
  Color graphics output    ==> NO      (YES/NO/MODIFY)
  Summarize by             ==> A1      (An-n/Dn-n/TABLE)
  Data Set Group Table     ==> _____ (If TABLE specified)
  Summary level           ==> ZONE    (DATE/HOUR/ZONE)

DOWN - Select  ENTER - Validate END - Run
-----

=====

----- StorageMate Data Filtering -----
Command ==>

Element Oper Value(s)
ENDTS      EQ  05APR91
ZONE       EQ  1
IACCT      EQ  _____

UP - Select    ENTER - Validate
-----

```

Figure D-9. Execution Options Panel

Input from DBID 'A' was requested, because that unit contains the ASTEX data. The administrator wants to monitor activity for Friday, April 5, 1991. Unsure which cycle contains the data for that date, the administrator requests cycles 01 through 03, and specifies that only date 05APR91 be selected. This will cause CA MICS StorageMate to do the work of finding the data. As the critical time period for our application is defined by zone 1, only that zone is requested. The output should consist of a printer graphic report, and summarization should be done using the contents of the first accounting field. Summarization by zone should also be done.

The example shown above actually represents portions of two different panels. Movement between the main report options panel, the data filtering panel, and the selection panel described below is performed via the ISPF UP and DOWN keys.

Because this facility allows reporting of multiple indicators, the administrator now uses the DOWN key to request display of the element selection panel. This causes the panel shown in Figure D-10 to be displayed.

```

----- ASTEX / Top Resource Consumers -----
Command ==>

Enter A-Z to select elements, then press UP to return.
I/O Activity
  _ Total I/O Duration                _ Total Number of I/O Requests
  _ Total IOS Queue Time              _ Total RPS Time
  _ Total I/O Pending Time            _ Total Seek I/O Requests
  _ Total I/O Disconnect Time          _ Average Seek Distance (Cylinders)
  _ Total I/O Connect Time             _ Average Response Time in MS

Cache Activity
  X Total Cache Effective (Hit) Count _ Total Tracks Loaded
  _ Total Cache-Candidate I/O Requests _ Total DFW Tracks Loaded
  _ Total Non-Candidate I/O Requests   _ Total NVS Overflow Count

Exceptions/Errors
  _ Total Dispatching Exceptions        _ Total Volume Exceptions
  _ Total Dsname Exceptions             _ Total Cross-Volume Exceptions
  _ Total Path Exceptions               _ Total I/O When Volume Reserved
  _ Total Seek Exceptions

UP - Report    DOWN - Filter  ENTER - Validate
-----

```

Figure D-10. Reporting Element Selection Panel

In this case, the administrator is only interested in selecting one element, although multiple elements could be selected for this single execution. The element 'Total I/O Requests Cached' has been selected, because that will report the number of I/O requests that were handled by cache, rather than by having to access DASD directly. Once this selection is made, the administrator presses UP to return to the previous panel, and then END to generate the report.

The resulting report produced from these two panels appears in Figure D-11.




Figure D-11. Sample Report Output

The results from the report are as the administrator expected, except that 5.25% of the cache activity is shown for data sets belonging to the Payroll department. That is not a critical application, and should not be cached.

To investigate this further, the administrator will run this report again, this time making some minor changes as shown on the panel in Figure D-12.

```

----- ASTEX / Top Resource Consumers -----
Command ==>

CA MICS Input Definition:
  Unit DBID(s) ==> A          Cycles ==> 01 - 03

Report Options:
  Produce reports          ==> DETAIL   (CHART/LIST/BOTH or DETAIL)
  Color graphics output    ==> NO      (YES/NO/MODIFY)
  Summarize by             ==> A1       (An-n/Dn-n/TABLE)
  Data Set Group Table     ==> _____ (If TABLE specified)
  Summary level            ==> ZONE     (DATE/HOUR/ZONE)

DOWN - Select  ENTER - Validate END - Run
-----

=====

----- StorageMate Data Filtering -----
Command ==>

Element Oper Value(s)
ENDTS      EQ  05APR91
ZONE       EQ   1
IACCT      EQ  PAYROLL

UP - Select    ENTER - Validate
-----

```

Figure D-12. Execution Options Panel

For this execution, only two changes have been made. An additional data filtering statement was coded to test the first accounting field for a value of PAYROLL, so that only the data from the payroll department will be selected. Secondly, a Detail report was requested instead of a Chart. This will give more detail as to how the values on the first report were derived. It is also not necessary to use the secondary element selection panel this time, because the elements from the last report will be remembered and used. Thus, the administrator simply uses the END key to execute the report.

The report produced by these options is shown in Figure D-13.

ASTEX / TOP RESOURCE CONSUMERS (STGEPF) CA MICS StorageMate - DETAIL XYZ PHONE SALES			
----- System Identifier=SYSA Owner=PAYROLL Date=05APR91 Zone=1 -----			
Time	Total Cache Effective (Hit) Count	Data Set Name	Volume Serial Number
9:30	11050	NWEST.PAYROLL.TST.FILE	CRIT01
9:30	0	CENT.PNAME.FILE	TS0023
10:00	0	AGENT5.ISPPROF	TS0055
10:00	26465	NWEST.PAYROLL.TST.FILE	CRIT01
10:00	0	AGENT5.ISPPROF	TS0055
10:30	22784	NWEST.PAYROLL.TST.FILE	CRIT01
10:30	0	CENT.PNAME.FILE	TS0023
11:00	10345	NWEST.PAYROLL.TST.FILE	CRIT01
11:00	0	CENT.PAYRATES.FILE	TS0061
11:30	5899	NWEST.PAYROLL.TST.FILE	CRIT01
-----	-----		
ZONE	76543		
	=====		
	76543		

Figure D-13. Sample Report Output

As expected, some of the data sets belonging to payroll have been accessed but had no cache activity. But one data set, NWEST.PAYROLL.TST.FILE, seems to have been responsible for all the cache activity during this measurement interval. The Storage Administrator notices immediately that this data set resides on volume CRIT01, which is one of the critical volumes under the cache controller.

The only task that remains is for the Storage Administrator to move the identified data set to a non-cached volume. This should reduce cache activity by approximately 5%, if this was a representative measurement interval.

Appendix E: NON-VSAM DATA SET ACTION-RELATED INTERFACE

While information about actions such as OPEN, CLOSE, DELETE, and RENAME performed on your non-VSAM data sets does not currently reside in CA MICS, StorageMate provides an external interface that allows you to place data set action-related information in a non-CA MICS SAS data set, and read that data set with CA MICS StorageMate. This section documents the steps you must take to create this external SAS data set before attempting to use this information in the StorageMate (STGEJD) report. You need not be concerned about doing the items described here unless you plan to use non-VSAM data set actions as an input source to StorageMate.

A sample job is distributed with this product as an example of how non-VSAM data set action information can be placed into a SAS data set and then read as input by CA MICS StorageMate. This job is contained in the following member:

o sharedprefix.MICS.CNTL(STGEJD)

This sample is designed to use input from your SMF data sets as input into the CA MICS StorageMate product. This job will process SMF record types 14, 15, 17, and 18, based on the parameters set in the MACRO call for %CREATEDB. You can select the following non-VSAM data set actions to report on OPEN, CLOSE, DELETE, and RENAME. Simply code the MACRO call as follows:

```

/*****
/* CODE THE FOLLOWING MACRO AS FOLLOWS:
/*
/* ACTION SMF RECORD
/* -----
/* OPEN 14
/* CLOSE 15
/* DELETE 17
/* RENAME 18
/* ALL 14,15,17,18
/*
/* %CREATEDB([OPEN][CLOSE][DELETE][RENAME][ALL])
/*
/* EXAMPLES: %CREATEDB(OPEN DELETE);
/* %CREATEDB(OPEN CLOSE);
/* %CREATEDB(OPEN CLOSE DELETE RENAME);
/* %CREATEDB(ALL);
*****/

```

Usage Notes

A. SMF Record Creation

The SMF 14 record is created for non-VSAM direct access data sets, tape data sets, VIO data sets defined by DD statements, or dynamically allocated and opened for RDBACK or INPUT processing. The SMF 15 record is created for non-VSAM data sets when they are opened for OUTPUT, UPDAT, INOUT, or OUTIN processing.

The SMF 17 record is created when a non-temporary or temporary non-VSAM data set is scratched. The SMF 18 record is created when a non-VSAM data set defined by a DD statement is either explicitly or implicitly renamed. CA MICS creates one observation for each SMF record processed.

B. Data Volume Issue for SMF Record Type 14

Be selective in your data collection for OPEN actions performed on your data sets. The data volume is very high for this record type. Remember, a record is created for every non-VSAM data set that is opened for OUTPUT, UPDAT, INOUT, or OUTIN processing.

C. Concatenated Sequential data sets

Records created for this type of data set will contain only the DDNAME for the first data set in the concatenation. All other data set DDNAMEs in the concatenation will be blank.

D. EXCP Counts

Use caution when relying on the EXCP counts from the SMF 14 and 15 records, because they may be invalid. The EXCP counts will only be valid if there is one occurrence of the DDNAME in the entire job. When there are multiple occurrences of the DDNAME within the job, the EXCP counts are accumulated across multiple opens within the same step. Due to this accumulation, EXCP counts are invalid.

E. Multiple Volume data sets

With SMF record types 17 and 18 representing the DELETE and RENAME data set actions, there is a recurring section in the record for data sets spanning multiple volumes. CA MICS StorageMate will store up to five volume serial numbers in elements VOLUME01 through VOLUME05. Element VOLSER will be set to the first volume (VOLUME01).

F. Concatenated BPAM data sets

If no member name is provided in the JCL for a PDS data set, the access method is BPAM. If this is the access method for the data set then only one 14 or 15 SMF record is generated. The record will contain only the DSNAME of the first data set in the concatenation. Although the record contains a UCB for each concatenation CA MICS StorageMate will create one observation for the entire record.

G. ISAM data sets

The ISAM sections of SMF record types 14 and 15 are not processed.

Appendix F: VSAM DATASET ACTIVITY-RELATED INTERFACE

While information about VSAM dataset action-related information and Integrated Catalog Facility activity does not currently reside in CA MICS, StorageMate provides an external interface that allows you to place this information in a non-CA MICS SAS dataset, and then read that dataset with CA MICS StorageMate. This section documents the steps you must take to create this external SAS dataset before attempting to use this information in StorageMate (STGEJF,STGEJG) reports. You need not be concerned about performing the items described here unless you plan to use VSAM dataset actions or ICF activity as an input source to StorageMate.

A sample job is distributed with this product as an example of how VSAM dataset action and Integrated Catalog Facility activity information can be placed into a SAS dataset and then read as input by CA MICS StorageMate. This job is contained in the following member:

- o sharedprefix.MICS.CNTL(STGEJFG)

This sample is designed to use input from your SMF data sets as input to the CA MICS StorageMate product. This job will process SMF record types 61, 62, 64, 65, and 66 based on the parameters set in the MACRO call for %CREATEDB.

You have an option to collect Integrated Catalog Facility activity information from the SMF 61, 65, and 66 records by specifying the ICFCAT parameter. This will create a SAS file called SMF6156 containing one observation for every record which was INSERTED, UPDATED, or DELETED from an ICF controlled catalog. You also have the option to collect VSAM dataset activity information from the SMF 62 and 64 records by specifying the OPENCLOSE parameter. This will produce two SAS files, SMF62 and SMF64. The SMF62 file will contain OPEN requests for VSAM datasets, and SMF64 will contain CLOSE requests for VSAM datasets.

Simply code the MACRO call as follows:

```
/******  
/* CODE THE FOLLOWING MACRO AS FOLLOWS:  
/*  
/*      PARAMETER      SMF RECORD      ACTION  
/*      -----      -  
/*      ICFCAT      61,65,66      INS. UPD. DEL.  
/*      OPENCLOSE      62,64      OPEN, CLOSE  
/*      ALL      61,62,64,65,66  
/*  
/* %CREATEDB([OPENCLOSE][ICFCAT][ALL])  
/*  
/* EXAMPLES: %CREATEDB(OPENCLOSE);  
/*           %CREATEDB(OPENCLOSE ICFCAT);  
/*           %CREATEDB(ICFCAT);  
/*           %CREATEDB(ALL);  
/******  
%CREATEDB(ALL);
```

Usage Notes

A. SMF Record Creation

VSAM Dataset Open and Close Activity

SMF record type 62 is written at the successful or unsuccessful opening of a VSAM component or cluster. This record identifies the VSAM component or cluster, gives the VSAM catalog in which the object is defined, and the volume on which the catalog and object are stored. It identifies the job by job log identification and user identification.

SMF record type 64 is written when a VSAM component or cluster is closed, VSAM must switch to another volume to continue to read or write, or there is no more space available for VSAM to continue processing. If a cluster is closed, one record is written for each component in the cluster.

VSAM Integrated Catalog Facility Activity

SMF record type 61 is written during any processing that results in a DEFINE request to Catalog Management Services. One record is written for each record inserted or updated in a catalog. This record identifies the entry being defined (the catalog in which the catalog record is written) and identifies the job by job log and user identifiers.

SMF record 65 is written during any processing that results in a DELETE request to Catalog management services. One type 65 record is written for each record updated or deleted from a catalog. Record type 65 identifies the entry being deleted and the catalog in which the catalog record is updated or deleted. It indicates whether a VSAM cluster or non-VSAM dataset was scratched, or only catalog information was deleted. It also indicates the job by job log identification and user identification.

Record type 66 is written during any processing that results in an ALTER request to Catalog Management Services. One type 66 record is written for each record written or deleted from a catalog. Record type 66 identifies the entry being altered and the catalog in which the catalog record is written or deleted. It indicates if the entry was renamed and, if so, gives the old and new names of the entry. It also indicates the job by job log identification and user identification.