

CA MICS® Resource Management

Performance Manager 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 Product Installation	9
1.1.1 How to Enable Neugents Technology in CA MICS.....	10
1.1.2 Foreground Execution of SAS.....	12
1.2 MVS Behavior Tracking and Analysis.....	13
1.3 I/O Analysis and Tape Drive Availability	15
1.4 Online Facility Tutorial	16
1.4.1 Introduction to the Online Facility	17
1.4.2 Getting Started with the Online Facility	21
1.4.3 Performance Manager Analysis Inquiries	41
1.4.4 Replaying Cataloged Output	66
 Chapter 2: PERFORMANCE REPORTING ANALYSIS	 67
2.1 ILM Analysis.....	68
2.1.1 Functional Description	69
2.1.2 Usage Guidelines.....	70
2.1.3 Standard Output	75
2.1.4 ILM Concepts.....	86
2.1.5 Requirements.....	87
2.2 Real Storage Management Analysis	89
2.2.1 Functional Description	90
2.2.2 Usage Guidelines.....	90
2.2.3 Standard Output	91
2.2.4 MVS Concepts	130
2.2.5 Analytic Tutorial	184
2.2.6 Case Study	184
2.2.7 Component Operation	185
2.2.8 Requirements.....	194
2.3 Auxiliary Storage Management Analysis	196
2.3.1 Functional Description	198
2.3.2 Usage Guidelines.....	200
2.3.3 Standard Output	204
2.3.4 MVS Concepts	215
2.3.5 Case Study	230
2.3.6 Component Operation	242
2.3.7 Requirements.....	250

2.4 I/O Configuration Analysis.....	251
2.4.1 Functional Description	252
2.4.2 Usage Guidelines.....	259
2.4.3 Shared DASD Analysis Inquiry	260
2.4.4 DASD I/O Component Analysis Inquiry	280
2.4.5 DASD Skew Analysis Inquiry	307
2.4.6 DASD Multiple Volume Analysis Inquiry	328
2.4.7 Tape Drive Simulation Analysis Inquiry	356
2.4.8 FICON Channel Performance Analysis Inquiry	424
2.5 Virtual Storage Management Analysis	449
2.5.1 Functional Description	450
2.5.2 Usage Guidelines.....	450
2.5.3 Standard Output	455
2.5.4 Virtual Storage Concepts in MVS/XA	483
2.5.5 Case Study	501
2.5.6 Component Operation	511
2.5.7 Requirements.....	523
2.6 PR/SM LPAR Performance Analysis	524
2.6.1 Functional Description	525
2.6.2 PR/SM LPAR Concepts.....	526
2.6.3 Standard Output	529
2.6.4 Usage Guidelines.....	542
2.6.5 Reserved.....	546
2.6.6 PR/SM LPAR Analysis Operation	546
2.6.7 Requirements.....	555
2.7 Graphic Analysis	556
2.7.1 CPU Graphs	557
2.7.2 Storage Graphs.....	566
2.7.3 TSO Service Graphs	571
2.7.4 I/O Tracking Graphs	575
2.7.5 Paging Graphs	580
2.7.6 Producing Graphs.....	589
2.8 Sysplex Analysis	613
2.8.1 Functional Description	615
2.8.2 Usage Guidelines.....	616
2.8.3 Standard Output	619
2.8.4 Parallel Sysplex Concepts	637
2.8.5 Component Operation	648
2.9 WLM Goal Mode Analysis	664
2.9.1 Functional Description	665
2.9.2 Usage Guidelines.....	667
2.9.3 Standard Output	673

2.10 Product Measured Usage Analysis	700
2.10.1 Functional Description	701
2.10.2 Usage Guidelines	702
2.10.3 Standard Output	708
2.10.4 Subcapacity Pricing Concepts	722
2.10.5 Reporting Requirements	722

Chapter 3: PERFORMANCE ANALYSIS TOOLS **727**

3.1 Batch Initiator Simulation	727
3.1.1 Functional Description	728
3.1.2 Usage Guidelines	729
3.1.3 Standard Output	730
3.1.4 Analytic Technique Tutorial	739
3.1.5 Case Study	741
3.1.6 Component Operation	754
3.1.7 Requirements	780
3.2 Relative Importance	780
3.2.1 Functional Description	781
3.2.2 Usage Guidelines	782
3.2.3 Standard Output	783
3.2.4 Component Operation	784
3.2.5 Case Study	793
3.2.6 Requirements	797
3.3 Data Clustering Analysis	798
3.3.1 Functional Description	800
3.3.2 Usage Guidelines	801
3.3.3 Standard Output	811
3.3.4 Analytic Technique Tutorial	833
3.3.5 Component Operation	845
3.3.6 Case Studies	866
3.3.7 Requirements	883

Appendix A: MESSAGES **885**

Messages	885

Chapter 1: OVERVIEW

The goal of performance management is to enable the most effective use of the current computing system to meet the goals of the organization. This consists of setting goals for system performance, arranging hardware and software to achieve the desired performance, tracking/analyzing actual achieved performance, and repeating the cycle as needed to adapt to changes in workload, configuration, policy, or newly observed system behavior.

The CA MICS Performance Manager Option gives you a set of tutorials and software tools to facilitate the performance management process. This component, together with the CA MICS database, assists you in tracking and analyzing system performance.

This section contains the following topics:

[1.1 Product Installation](#) (see page 9)

[1.2 MVS Behavior Tracking and Analysis](#) (see page 13)

[1.3 I/O Analysis and Tape Drive Availability](#) (see page 15)

[1.4 Online Facility Tutorial](#) (see page 16)

1.1 Product Installation

Instructions for installing the CA MICS Performance Manager Option were included with your CA MICS product distribution tape. In the second step of these instructions you ran the load job to copy all required product modules and members into your CA MICS libraries.

The following sections detail additional requirements and considerations:

- 1 - How to Enable Neugents technology in CA MICS
- 2 - Foreground Execution of SAS

1.1.1 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

1.1.1.1 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.

1.1.1.2 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.

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.

1.1.2 Foreground Execution of SAS

Some CA MICS Performance Manager Option inquiries executing in the foreground environment require that SAS be executed via the SASCP command rather than by using the CALL command. These inquiries create temporary PROC FORMATS in ddname LIBRARY during their execution. SASCP searches LIBRARY for temporary formats if the format is not found in ddname SASLIB. SAS's CALL mode, however, builds a directory of existing formats within SASLIB files during its start-up processing and therefore will never find any temporary formats created afterwards during the inquiry processing.

The CA MICS Administrator specifies the SAS foreground interface mode along with other SAS Global parameters through the CA MICS Administrator Facility options. Refer to Section 4.4.2 of the CA MICS Planning, Installation, Operation, and Maintenance Guide (PIOM) for more information.

1.2 MVS Behavior Tracking and Analysis

The CA MICS Performance Manager Option includes tutorials and reports for tracking MVS by analyzing data about the System Resource Manager (SRM), Real Storage Manager (RSM), Virtual Storage Manager (VSM), and the Auxiliary Storage Manager (ASM). This is described briefly in this section and more fully in Sections 2.2, and 2.3. Additionally, with the introduction of MVS/ESA SP 5.1 and Parallel Sysplexes, the Performance Manager incorporates reports for analyzing data about the functioning of Sysplex configurations (SPA). Section 2.8 describes these reports in more detail.

The Installation Performance Specifications (IPS) control the SRM, which is a component of the MVS control program. The SRM controls the distribution of system resources (CPU, I/O, and real storage) by controlling the amount of time each address space is allowed to execute. SRM's objective is to achieve optimal use of system throughput. Details on IPS analysis are no longer provided in this guide.

In controlling these resources, the SRM must interface with the other major MVS components, including RSM, VSM, and ASM.

RSM maintains the control blocks that translate real and virtual storage addresses, makes the determination of working set contents when an address space is swapped, and performs the page replacement operations. RSM analysis is discussed in detail in Section 2.2.

ASM manages the page and swap data sets that comprise the virtual space. The page data sets provide the auxiliary storage that backs up the storage frames. Section 2.3 describes the operation of ASM, detailing its real storage control mechanisms and their relation to the other system components. It also discusses the algorithms controlling the allocation of the auxiliary storage devices and provides guidelines for tuning an ASM configuration.

MVS/ESA SP 5.1 provides support for the System 390 Parallel Sysplex, which consists of hardware, software, and optional services. The Parallel Sysplex supports full data sharing among the processors in the sysplex and dynamic workload balancing across the system. Section 2.8 describes the operation of a sysplex, including use of the Coupling Facility, which provides the means for multisystem data sharing by authorized subsystems in a Parallel Sysplex. It also discusses the analysis features of the CA MICS MVS Performance Manager that assist in managing and tuning a sysplex environment.

1.3 I/O Analysis and Tape Drive Availability

The configuration and utilization of input/output devices is a critical aspect of system performance management. To support this part of system performance management, the CA MICS Performance Manager Option enables configuration analysis, including channel/path/device utilization analysis, DASD data set placement contention reporting, and tape drive availability analysis.

The I/O Analysis software for DASD implements four different algorithms: the Shared DASD Analysis Algorithm, the I/O Component Analysis Algorithm, the DASD Skew Analysis Algorithm, and the Multiple Volume Mapping Algorithm.

Shared DASD allows multiple systems to be connected to a single volume. This feature is highly desirable for some applications' volumes because critical applications can be moved to another system in the event of a processor failure. This feature also allows a single copy of critical data to be shared by multiple systems.

Shared DASD often introduces significant performance problems, however, by allowing workloads on one system to adversely affect activities on other systems. Because RMF reports on the I/O activities of individual systems rather than the I/O activity of an installation as a whole, it is necessary to summarize RMF data from multiple systems to evaluate the performance of a shared device.

The I/O service time is the most significant part of response time for all but the most CPU bound applications. As CPU internal processing speeds increase, performance problems that result from I/O delays increase.

DASD skew is a condition that existed prior to the introduction of controllers with RPS avoidance, such as EMC and 3990-6 controllers. It stems from the nature of rotational position sensing (RPS) devices and RPS miss probabilities. If you are not using a controller with RPS avoidance, DASD skew occurs when one or two of the devices on a string are heavily used and cause performance problems for the remaining devices on the string.

Dual density DASD devices or devices that offer adequate storage capacities offer cost and floor space reductions. The conversion to this type of device must be well planned. Devices that are to be mapped together to form a dual density volume should not be experiencing high I/O activity during

the same time periods. If such volumes are combined, severe arm contention (that is, delays) will result.

The Tape Drive Analysis facility allows you to analyze time periods that have significant tape drive availability problems, and to model the effects of moving existing work to different time periods, as well as the effects of increasing tape usage. You can investigate both the effects on performance of rescheduling existing jobs and the capacity implications of increased tape drive usage.

See Section 2.4 I/O Configuration Analysis, for details on all these algorithms and the reports that are produced.

1.4 Online Facility Tutorial

The online facility provides the means for using the CA MICS Performance Manager Option. This section introduces the online facility and explains how to use it. Additional online tutorials provide extensive guidance in using the features of the CA MICS Performance Manager Option.

If you are familiar with MICF (CA MICS Information Center Facility), you may notice many similarities between it and the CA MICS Performance Manager Option's online facility. The online facility uses many of the MICF features. Any experience you have with MICF should benefit you as you begin using the CA MICS Performance Manager Option online facility.

The remainder of this guide contains detailed descriptions of each of the CA MICS Performance Manager Option features and their operation using the online facility.

This section discusses the following topics:

- 1 - Introduction to the Online Facility
- 2 - Getting Started with the Online Facility
- 3 - Performance Management Analysis Inquiries
- 4 - Replaying Cataloged Output

1.4.1 Introduction to the Online Facility

The following sections provide an overview of the purpose and major features of the online facility:

- 1 - Overview of the Online Facility
- 2 - Major Features of the Online Facility

1.4.1.1 Overview of the Online Facility

The online facility provides a convenient way for you to use the CA MICS Performance Manager Option to rapidly obtain the results you need. With the online facility, you do not need an in-depth knowledge of CA MICS or SAS to effectively use the CA MICS Performance Manager Option for building, maintaining, and analyzing performance management database files. The online facility's automatic validation procedures prevent you from making invalid requests, such as trying to use data elements that do not exist in the CA MICS files. Automatic validation also prevents syntax errors during subsequent analysis of performance management database files.

Following MICF's format, the online facility allows you to create programs called analysis inquiries. You can design these inquiries to analyze that data and produce tabular and graphics reports.

Online screens are provided for each reporting feature of the CA MICS Performance Manager Option. These panels and the chapters in this guide that discuss them lead you through the use of each of these features.

A comprehensive, interactive tutorial facility assists you in learning and using the online facility. In addition, you have online access to all CA MICS documentation while you are using the online facility.

The following steps briefly describe how the online facility works:

- o You define the input, selection criteria, analysis, and report options.
- o Transparent to you, the online facility dynamically builds the analysis inquiry program from your specifications.
- o The inquiry is executed interactively or in batch, and performs an analysis upon the data in your database files. The reports and diagnostic output produced during the inquiry execution are displayed at your terminal or routed to a printer.

1.4.1.2 Major Features of the Online Facility

The online facility is a menu-based system that operates under IBM's Interactive Systems Productivity Facility (ISPF) and makes use of many of the features of the CA MICS Information Center Facility (MICF). If you are familiar with ISPF commands, you will find that the online facility menus behave like ISPF menus. The major features of the online facility include the following:

- o Inquiry composition
- o Inquiry modification
- o Execution-time specifications
- o Output replay

Each of these features, described below, is generally applicable to both file definition inquiries and analysis inquiries.

INQUIRY COMPOSITION

By completing one or a few online screens, you can create any number of performance management inquiries. These inquiries allow you to analyze specific areas of MVS (such as real and auxiliary storage and the DASD subsystem), to characterize workloads via clustering analysis, and to simulate the effects of initiator restructuring or workload demands on tape drive availability.

The performance management inquiry screens are designed to be easy to use. You need to specify only those fields required for the operation of a function. Many fields have default values defined to make operation even easier for you.

One of the features of the performance management screens is that your specifications for various fields are saved by the system so that you need not reenter them unless you want to change the field value. These fields are preserved not only as you go from function to function, but they are also preserved from day to day as you execute the performance management software.

The fields that retain your previous specifications are:

- o Timespan
- o Cycle(s)
- o Database IDs
- o Allow runtime overrides
- o Report by hour

- o Generate report by Zone
- o Date selection
- o Hour selection
- o Zone selection
- o SYSID selection
- o Performance Group selection
- o Performance Domain selection
- o Performance Period selection
- o Program selection
- o Job selection

INQUIRY MODIFICATION

Once created, an inquiry can be stored in your private catalog, recalled, copied, deleted, or modified. You make changes to an inquiry using the same online facility screens that you originally used to create the inquiry. You do not, however, need to redisplay screens that did not change.

EXECUTION-TIME SPECIFICATIONS

You can define analysis inquiries so that key inquiry parameters can be specified at the time the analysis inquiry is executed. This facility allows you to report against varying data without changing the analysis inquiry.

OUTPUT REPLAY

The online facility allows you to save the output produced when you execute a file definition or analysis inquiry, and replay it later. The output is saved by inquiry name, date, and time. Your CA MICS System Administrator can place inquiry outputs into a shared replay catalog that can be used throughout your installation.

When you replay an inquiry output, you can either view it on your terminal or send it to a hardcopy device. If you are replaying color graphics output, you do not need to use the type of device that was specified at the time the inquiry was originally executed. You can specify the color terminal, color printer, or color plotter type to use at the time you request the replay.

1.4.2 Getting Started with the Online Facility

After you read this section, you will be able to start the online facility and know how to use its menus, commands, and help screens. In addition, you will be able to customize the online facility to better meet your needs.

The online facility contains extensive online tutorials (help screens) that you can request at any point in an online facility session. These screens provide detailed information about the operation you are specifying at the time you request help.

When you leave the tutorials, the online facility returns you to the place in your work session from which you requested help. Ask for help by typing HELP after the Command prompt (or the Option prompt on some screens) at the top of the screen. Your PF1 key is usually set to do this for you. When you request help in this manner, you are immediately shown a tutorial that is relevant to the screen on which you are working. To leave a tutorial, type END after the Command (or Option) prompt at the top of your screen. Your PF3 key is usually set to do this for you.

This remainder of this section discusses the following topics:

- 1 - Operating Conventions of the Online Facility
- 2 - Starting the Online Facility
- 3 - Customizing the Online Facility

1.4.2.1 Operating Conventions of the Online Facility

The CA MICS Performance Manager Option online facility follows the same format and operating conventions as the CA MICS Information Center Facility (MICF).

Both the online facility and ISPF are very easy to use. Even if you have very little data processing experience, you will be able to learn the basics very quickly.

1.4.2.2 Starting the Online Facility

Start the online facility by selecting it from the CA MICS Workstation Facility (MWF) menu. The method of reaching the CA MICS Workstation Facility varies from installation to installation.

There are two common ways to start MWF. If neither of these works, ask your CA MICS System Administrator to tell you how to start MWF. This section discusses the following topics:

- 1 - Selecting MWF from an ISPF Menu
- 2 - Starting MWF by Executing a CLIST

1.4.2.2.1 Selecting MWF from an ISPF Menu

Your CA MICS System Administrator may have installed CA MICS so that you can reach the CA MICS Workstation Facility by selecting an option on your installation's ISPF primary option menu or another ISPF menu. If this is the case, when you start ISPF you will see a menu similar to the one shown in Figure 1-1.

```

/----- ISPF/PDF PRIMARY OPTION MENU -----\
|OPTION ==> MWF|
|
|                                USERID - yourid |
| 0 ISPF PARMs - Specify terminal and user parameters  TIME - 10:19 |
| 1 BROWSE - Display source data or output listings  TERMINAL - 3278 |
| 2 EDIT - Create or change source data  PF KEYS - 24 |
| 3 UTILITIES - Perform utility functions |
| 4 FOREGROUND - Invoke language processors in foreground |
| 5 BATCH - Submit job for language processing |
| 6 COMMAND - Enter TSO command or CLIST |
| 7 DIALOG TEST - Perform dialog testing |
| 8 LM UTILITIES- Perform library management utility functions |
| C CHANGES - Display summary of changes for this release |
| MWF CA MICS - CA MICS Workstation Facility (MWF) |
| T TUTORIAL - Display information about ISPF/PDF |
| X EXIT - Terminate ISPF using log and list defaults |
|
|Enter END command to terminate ISPF.
|
|
|
|
\-----/

```

Figure 1-1. Selecting MWF from the ISPF/PDF Primary Option Menu

Select the CA MICS Workstation Facility option by typing MMF after the option prompt, and pressing ENTER. You will see either the Welcome New CA MICS User panel (Figure 1-2) or the primary CA MICS Workstation Facility (MMF) panel (Figure 1-3). The Welcome New CA MICS User panel appears the first time you use the online facility so that a set of user tables can be allocated for you. After they are initially allocated, you will go directly to the MMF panel in future start-ups.

```

/----- Welcome New CA MICS User -----
|Command ==>;
|
| Allocation failed for:  youruserid.MICS.USER.TABLES
|
| Press ENTER to automatically allocate this data set and continue.
| Press END to terminate without allocating the data set.
| Note: Allocation failure is normal for new users.  If you were an
| existing user, you may want to terminate and investigate the failure.
|
| The following parameters are required.  Specify sufficient space for
| your MICS parameters and inquiries, your documentation bookmarks, etc.
|
| Space Type          ==> CYLINDERS (CYLINDERS/TRACKS/BLOCKS)
| Primary Space       ==> 2                      Secondary Space ==> 1
| Directory Blocks    ==> 25                     Block Size          ==> 6160
|
| The following parameters may be optional in your environment.  However,
| you should specify a Unit/Volume combination that will not be subject
| to scratching.
|
| Unit Name           ==> (3380,3350,DISK,SYSDA,etc.)
| Volume Serial       ==> _____ (Site dependent value)
|
\-\-\-----
/

```

Figure 1-2. Welcome New CA MICS User Panel

Your CA MICS System Administrator may have specified a default unit name or volume serial identification. If so, press ENTER to allocate the data set for your user tables. Depending on your site's requirements for allocating permanent data sets for TSO users, the Unit Name and Volume Serial parameters will be either optional or required. If you need to change the Unit Name or Volume Serial parameters, enter the correct information after the prompts and press ENTER. MWF allocates the data set for your user tables and displays the primary MWF panel shown in Figure 1-3.

```
----- 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,
Performance Manager and StorageMate.
- 5 - CA MICS Administrator Facility (MAF)
Operational status and tracking, installation, modification,
maintenance, and authorization.
- T - Tutorial
- X - Exit

```
. . . . .
```

Figure 1-3. CA MICS Workstation Facility (MWF) Panel

Select Option 4, Management Support Applications, from the CA MICS Workstation Facility (MWF) panel. The Management Support Applications panel (Figure 1-4) is displayed.

----- 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 - Performance Manager
Analyze performance of MVS systems.
- 5 - StorageMate
Report and analyze storage resources.

.

Figure 1-4. Management Support Applications

Select Option 4, Performance Manager, from the Management Support Applications panel. The Performance Manager selection panel (Figure 1-5) is displayed. This panel allows you to define a new inquiry or select an inquiry from the Shared catalog or your Private catalog. You can optionally modify your MICF options or review reports and graphics generated by previously executed inquiries.

```

/----- Performance Manager -----\
|Option ==>
|
|    0 - MICF Options
|        Review/modify MICF parameters and default options.
|
|    1 - Performance Reporting
|        Compose/execute a new private performance inquiry.
|
|    2 - Performance Analysis Tools
|        Explore performance data characteristics and simulate
|        initiator structures.
|
|    3 - Performance Manager Catalog
|        Modify/execute previously built inquiries saved
|        in either Shared or Private Performance catalogs.
|
|    4 - Reports and Graphics
|        Review printed reports and color graphics output
|        from previously executed inquiries.
|
|    T - Tutorial
|
|    X - Exit
|
\-\-\-\-----/

```

Figure 1-5. Performance Manager Selection Panel

Select option 1 to display the Performance Reporting panel (Figure 1-6). Select Option 2 to display the Performance Analysis Tools panel (Figure 1-7). If you would like to view previously saved performance queries, or wish to view the shared performance queries delivered with the product, select option 3 to display the Performance Manager Catalog panel (Figure 1-8).

```
/----- Performance Reporting -----\
|Option ==>
|
|   Inquiry Name   ==> ADHOC
|       Title     ==> Adhoc Inquiry - Will Not Be Saved
|       Author    ==> _____
|
|Catalog Group Name ==> PERFORM
|       Title      ==> Performance and Resource Consumption
|
|   1 - General Systems Analysis
|       Analyze the ILM, RSM, ASM, and VSM.
|   2 - Input/Output Configuration (I/O) Analysis
|       Analyze DASD or tape I/O configuration usage and performance.
|   3 - PR/SM LPAR Analysis
|       Analyze CEC and LPAR utilization, configuration and performance.
|   4 - Graphic Analysis
|       Track and manage service level and performance indicators.
|   5 - Sysplex Analysis
|       Analyze configuration and coupling facility usage/performance.
|\-\\-\\-\\-----/
```

Figure 1-6. Performance Reporting Panel

```
/----- Performance Analysis Tools -----\  
|Option ==>  
|  
|    Inquiry Name  ==> ADHOC  
|                Title ==> Adhoc Inquiry - Will Not Be Saved  
|                Author ==> _____  
|  
|Catalog Group Name ==> PERFORM  
|                Title ==> Performance and Resource Consumption  
|  
|    1 - Relative Importance Analysis  
|        Utilizes proprietary Neugents technology to understand  
|        and quantify the effect of a list of data elements against  
|        a client specified analysis data element.  
|    2 - Data Clustering Analysis  
|        Utilizes proprietary Neugents technology to explore pattern  
|        similarities, and differences in data classifications.  
|    3 - Batch Initiator Simulation  
|        Simulates changes to batch workload or initiator structure.  
|\-\\-\\-\\-----/
```

Figure 1-7. Performance Analysis Tools Panel

1.4.2.2.2 Starting MWF by Executing a CLIST

If you cannot start MWF from a standard ISPF menu, ask your CA MICS System Administrator for the name of the CA MICS Workstation Facility CLIST. For purposes of this discussion, we will assume that it is `sharedprefix.MICS.CLIST(MWF)`.

First, sign on to TSO. From the TSO READY prompt, type

```
EX 'sharedprefix.MICS.CLIST(MWF)'
```

If you are a new CA MICS user, you will see the Welcome New CA MICS User panel shown in Figure 1-2. Your CA MICS System Administrator may have specified a default Unit Name or Volume Serial. If so, press ENTER to allocate the data set for your user tables. Depending on your site's requirements for allocating permanent data sets for TSO users, the Unit Name and Volume Serial is either optional or required. If you need to change the Unit Name or Volume Serial, enter the correct information after the prompts and press ENTER. MWF allocates the data set for your user tables and displays the primary CA MICS Workstation Facility (MWF) panel, which is shown in Figure 1-9.

If this is not the first time you have used MICS, you will go directly to the primary CA MICS Workstation Facility (MWF) 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, Performance Manager and StorageMate. 5 - CA MICS Administrator Facility (MAF) Operational status and tracking, installation, modification, maintenance, and authorization. - T - Tutorial - X - Exit
.

Figure 1-9. CA MICS Workstation Facility (MWF) Panel When Invoked from CLIST

Select Option 4, Management Support Applications, from the CA MICS Workstation Facility (MWF) panel. The Management Support Applications panel, shown in Figure 1-10, is displayed.

----- 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 - Performance Manager
Analyze performance of MVS systems.
- 5 - StorageMate
Report and analyze storage resources.




Figure 1-10. Management Support Applications Panel

The Performance Manager panel (Figure 1-11) is displayed when you select Option 4, Performance Manager, from the Management Support Applications panel. This panel allows you to define a new inquiry or select an inquiry from the Shared catalog or your private catalog. You can optionally modify your MICF options or review reports and Graphics generated by previously executed inquiries.

Select option 1 to display the Performance Reporting panel (Figure 1-12). Select Option 2 to display the Performance Analysis Tools panel (Figure 1-13). If you would like to view previously saved performance queries or wish to view the shared performance queries delivered with the product select option 3 to display the Performance Manager Catalog panel (Figure 1-14).

```

/----- Performance Manager -----\
|Option ==>                               |
|                                           |
|      0 - MICF Options                     |
|      Review / modify MICF parameters and default options. |
|                                           |
|      1 - Performance Reporting             |
|      Compose / execute a new private performance inquiry. |
|                                           |
|      2 - Performance Analysis Tools        |
|      Explore performance data characteristics and simulate |
|      initiator structures.                 |
|                                           |
|      3 - Performance Manager Catalog       |
|      Modify / execute previously built inquiries saved    |
|      in either Shared or Private Performance catalogs.   |
|                                           |
|      4 - Reports and Graphics              |
|      Review printed reports and color graphics output    |
|      from previously executed inquiries.                 |
|                                           |
|      T - Tutorial                               |
|                                           |

```

[illegible]

Figure 1-11. Performance Manager Selection Panel

```

/----- Performance Reporting -----\
|Option ==>
|
|    Inquiry Name  ==> ADHOC_
|    Title        ==> Adhoc Inquiry - Will Not Be Saved
|    Author       ==> _____
|
|Catalog Group Name ==> PERFORM
|    Title        ==> Performance and Resource Consumption
|
|    1 - MVS Analysis
|        Analyze ILM, RSM, ASM, and VSM.
|    2 - Input/Output Configuration (I/O) Analysis
|        Analyze DASD or tape I/O configuration usage and performance.
|    3 - PR/SM LPAR Analysis
|        Analyze CEC and LPAR utilization, configuration and performance.
|    4 - Graphic Analysis
|        Track and manage service level and performance indicators.
|    5 - Sysplex Analysis
|        Analyze configuration and coupling facility usage/performance.
\-----\
/

```

Figure 1-12. Performance Reporting Panel

```

/----- Performance Analysis Tools -----\
|Option ==>
|
|    Inquiry Name  ==> ADHOC_
|    Title        ==> Adhoc Inquiry - Will Not Be Saved
|    Author       ==> _____
|
|Catalog Group Name ==> PERFORM
|    Title        ==> Performance and Resource Consumption
|
|    1 - Relative Importance Analysis
|        Utilizes proprietary Neugents technology to understand
|        and quantify the effect of a list of data elements against
|        a client specified analysis data element.
|    2 - Data Clustering Analysis

```

```

|           Utilizes proprietary Neugents technology to explore pattern |
|           similarities, and differences in data classifications.      |
|           3 - Batch Initiator Simulation                             |
|           Simulates changes to batch workload or initiator structure. |
\-\-----
/

```

Figure 1-13. Performance Analysis Tools Panel

```

/----- Performance Manager Catalog -----\
|Command ==>                                     Scroll ==> PAGE |
|
|Catalog    ==> PRIVATE (PRIVATE/SHARED) Catalog Group ==> PERFORM |
|Select Mode ==> MODIFY (EXECUTE/MODIFY) Name (subset display) ==>_____ |
|
|Line Cmds: S Select I Insert R Repeat C Copy D Delete              |
|
|Cmd  Name          Title          Date    Time  Type |
| -  - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - |
| _  WKLD1    Data Clustering          yy/mm/dd 13:24 PERCLU|
|*****|
|
|
|
|
|
|
|
|
|
|
\-\-----
/

```

Figure 1-14. Performance Manager Catalog Panel

1.4.2.3 Customizing the Online Facility

This section provides a general description of how to customize the default parameters governing your online facility sessions. Parameters that are used by the CA MICS Performance Manager Option online facility are shared with those used by MICF. Consequently, changes you make to these parameters are reflected when you use MICF, and vice-versa.

When you use the online facility, you must specify how to allocate data sets, what printer formats to use, what color graphics devices you are using, and what ISPF profile you want to use with the online facility. Defaults for each of the parameters in these categories were included with MICF when it was shipped to your installation. Your CA MICS System Administrator has further tailored these defaults for your site. In addition, the online facility often gives you the opportunity to revise parameters temporarily for a particular inquiry.

For these reasons, we recommend that you start using the online facility with the default options unchanged. As you become more familiar with the online facility, you may wish to reset some of these defaults. To do so, select Option 0, MICF Options, on the CA MICS Information Center Facility (MICF) menu, or Option 0, MICF Options, on the Performance Manager menu. The selection menu is shown in Figure 1-15.

Chapter 1: OVERVIEW 37

Because each parameter type is fully described in the online tutorials, we will not discuss them in detail here. We will, however, walk you through one type of parameter change to show you what to expect during this process.

Some parameters that you may want to eventually change are:

- o Data Set Allocation Parameters
 - Space for the work data set
- o User Profile Parameters
 - Default job card
 - Default execution mode (foreground or batch)
 - Whether you will be prompted for foreground or batch execution
 - Whether you will be given a chance to edit your JCL before you submit a job

Suppose you want to change the default job card that is used each time you submit a batch job from the online facility. Select Option 1 on the MICF Options panel and press ENTER.

The MICF Parameter Update panel is shown in Figure 1-16. The top line contains the familiar Command and Scroll prompts. These are followed by the title "Updating User Profile Parameters," and the prompt, "Subset Display On." This prompt refers to the display under the title "Parameter Title/Parameter Value."

```

/----- MICF Parameter Update -----\
| Command ==>                               Scroll ==> CSR |
| Updating User Profile Parameters           |
| Subset Display On ==> -----             |
|                                     |
|                                     |
|-----|
| Batch Execution - Display generated jobstream before submitting jobs (YES/NO)? |
| ==> YES                               |
|                                     |
| Caps. - Translate ALL screen entry "title/label" fields to uppercase (YES/NO)? |
| ==> NO                               |
|                                     |
| Caps. - Translate ALL screen entry "value" fields to uppercase (YES/NO)?       |
| ==> YES                               |
|                                     |
| Confirm Delete - Require confirmation to delete an inquiry step (YES/NO)?      |
| ==> YES                               |
|                                     |
| Cycle(s) - Specify the default file cycle(s) for inquiry execution (cc-cc)     |
| ==> -----                         |
|                                     |
\-----/

```

Figure 1-16. Updating User Profile Parameters

The display on the MICF Parameter Update panel is an alphabetical list of all modifiable user profile parameters. Figure 1-16 shows the first five parameters in the list.

The line following each parameter contains a prompt arrow and the current value for the parameter. If you want to change the value, simply type the value you want over the current value.

There are four job card parameters in the list that you can reach by scrolling through the list. Note that the information changes under Parameter Title/Parameter Value, but the information above the line does not. If you know the first letter or letters of the parameter you wish to change, you can reach it more quickly by specifying this information in the "Subset Display On" field at the top of the panel. For example, if you subset the display on "Job", the parameters shown in Figure 1-17 are displayed. Note that capitalization is important in subsetting. If you enter "JOB" or "job", no parameter will be found.

```
/----- MICF Parameter Update -----\
| Command ==>                               Scroll ==> CSR |
| Updating User Profile Parameters                |
| Subset Display On ==> Job                      |
|                                     |
|               Parameter Title/Parameter Value |
| -----|
| Job Card 1 - Default Job statement #1 for ALL MICF batch processing. |
| ==> //jobname  JOB (accounting information), |
|                                     |
| Job Card 2 - Default Job statement #2 for ALL MICF batch processing. |
| ==> // MSGLEVEL=(1,1),MSGCLASS=A,PRTY=5 |
|                                     |
| Job Card 3 - Default Job statement #3 for ALL MICF batch processing. |
| ==> // |
|                                     |
| Job Card 4 - Default Job statement #4 for ALL MICF batch processing. |
| ==> // |
|                                     |
| ***** BOTTOM OF DATA ***** |
|                                     |
| \- \- \- -----/
```


Figure 1-17. Subsetting the Display on Job

The information shown for the default job cards varies from installation to installation. Your CA MICS System Administrator can provide you with the appropriate defaults for your installation.

The information on these job cards is supplied automatically each time you submit a batch job from MICF or from the online facility. By default, you will be given the opportunity to edit the generated JCL prior to submitting a batch job, which also allows you to edit your job card at execution time.

If you want to change one or more parameters on your default job card, enter the new parameters after the arrow prompt on Job Card 1, Job Card 2, Job Card 3, or Job Card 4. When you END from this panel (usually by pressing your PF3 key), the information you changed is saved. If you make a mistake and do not want to save your changes, type CANCEL after the Command prompt.

The other parameter update options work in the same manner.

1.4.3 Performance Manager Analysis Inquiries

This section introduces the online facilities that are available for creating, modifying, deleting, and executing performance management analysis inquiries. Subsequent chapters of this guide provide direction on the parameters necessary to perform each particular type of analysis.

The following sections address these topics:

- 1 - Creating an Analysis Inquiry
- 2 - Modifying an Existing Analysis Inquiry
- 3 - Repeating an Existing Analysis Inquiry
- 4 - Deleting an Analysis Inquiry
- 5 - Executing an Analysis Inquiry
- 6 - Specifying Common Extended Options

1.4.3.1 Creating an Analysis Inquiry

The process of creating an analysis inquiry is straightforward. You tell the online facility that you are about to create an inquiry. Then you give the inquiry a name and assign it to a catalog group (a discussion on catalog groups follows). Next, you define the analysis technique you plan to use. At this point, you can choose to cancel, save, or execute the inquiry. If you decide to execute the inquiry before saving it, you will have an opportunity to cancel, save, or modify the inquiry after execution.

To create an analysis inquiry, select Option 4, MVS Performance Manager, from the Management Support Applications panel (Figure 1-4). The Performance Manager panel (Figure 1-5) is displayed. This panel allows you to define a new inquiry or select an inquiry from the Shared catalog or your Private catalog. You can optionally modify your MICF options or review reports and graphics generated by previously executed inquiries.

Select option 1 from the Performance Manager panel to display the Performance Reporting panel (Figure 1-18) for defining new inquiries. If you have not previously saved any analysis inquiries, the Performance Reporting panel is automatically displayed. Selecting option 3 on the Performance Manager panel displays the Performance Manager Catalog panel (Figure 1-19) for selecting a previously defined inquiry.

```
/----- Performance Reporting -----\  
|Option ==>|  
|  
|   Inquiry Name ==> _____|  
|       Title ==> _____|  
|       Author ==> _____|  
|  
|Catalog Group Name ==> PERFORM|  
|       Title ==> Performance and Resource Consumption|  
|  
|   1 - General Systems Analysis|  
|       Analyze ILM, RSM, ASM, and VSM.|  
|   2 - Input/Output Configuration (I/O) Analysis|  
|       Analyze DASD or tape I/O configuration usage and performance.|  
|   3 - PR/SM LPAR Analysis|  
|       Analyze CEC and LPAR utilization, configuration and performance.|  
|   4 - Graphic Analysis|  
|       Track and manage service level and performance indicators.|  
|   5 - Sysplex Analysis|  
|       Analyze configuration and coupling facility usage/performance.|  
\-----/
```

Figure 1-18. Performance Reporting Panel

On the Performance Reporting panel you can specify the Inquiry Name (one to six characters), Title, and Author. The Catalog Group Name and Title can be specified for an existing catalog or you can change the catalog group to which an inquiry belongs. If you want to change either the Catalog Group Name or the Title, you can do so here. If you want to change any information on this menu later, you must make a copy of the inquiry to do so.

An Inquiry Catalog is a group of logically related inquiries that you assign to a specific group name. This allows you to subdivide your inquiries into distinct groupings for easier reference at a later time.

As shown in Figure 1-18, there are five different types of performance management processes you can use to create an analysis inquiry.

Choosing Option 1 allows you to analyze the ILM, RSM, ASM, or VSM areas of MVS. Sections 2.1 through 2.3 and Section 2.5 of this guide discuss this option and its subsequent panels.

Choosing Option 2 allows you to analyze DASD or tape I/O configuration usage and performance. Section 2.4 of this guide discusses this option and its subsequent panels.

Choosing Option 3 allows you to analyze PR/SM LPAR performance. Section 2.6 of this guide discusses this option and its subsequent panels.

Choosing Option 4 allows you to use graphic analysis to track and manage a number of key service level and performance related indicators. Section 2.7 of this guide discusses this option and the various categories of graphs provided.

Choosing Option 5 allows you to monitor the activity levels, configuration, and performance of cross-system functions (XCF) and the Coupling Facility (CF). Section 2.8 of this guide discusses this option and its subsequent panels.

[illegible]

Figure 1-19. Performance Manager Catalog Panel

As a user of the online facility, you have a private catalog of analysis inquiries. Each catalog is divided into catalog groups. These groups contain the actual inquiries. When the Catalog Group field contains no specification, the display portion of the panel displays the names of the various catalog groups that have been defined.

When you select a catalog group, either by typing S (Select) in the Cmd column or by typing in a catalog group name after the Catalog group prompt, the display portion of the panel shows the performance management analysis inquiries contained in that catalog group. To see a list of all inquiries in your catalog, type an asterisk (*) after the Catalog Group prompt. To see a list of all inquiries in all catalog groups starting with the same sequence of characters, code the sequence followed by an asterisk after the Catalog Group prompt. For example,

```
Catalog Group =====> S*
```

shows you the inquiries in all catalog groups starting with the letter S.

If the list of inquiries is long, you can scroll through it using the standard ISPF DOWN and UP commands. You can subset the list of inquiries by entering the beginning character(s) of the names you are interested in after the Name prompt. Only inquiries with names starting with these characters are shown.

The Select Mode controls the action of the S (Select) line command and the SELECT primary command. The two possible values are EXECUTE or MODIFY; the setting defaults to the last value specified for this parameter. Specify MODIFY to modify an existing inquiry. Specify EXECUTE to run the inquiry.

The information about each inquiry is displayed on a single line. The information includes the inquiry name, the title, the date and time the inquiry was created, and what type of analysis is performed in the inquiry.

You can create a new analysis inquiry by typing an I (Insert) on any line in the Cmd column, and pulling up Performance Reporting panel (Figure 1-18).

1.4.3.2 Modifying an Existing Analysis Inquiry

You can change an existing analysis inquiry and replace the original from the Performance Management Catalog panel (shown in Figure 1-20). Specify MODIFY for Select Mode, type an S (Select) in the Cmd column corresponding to the inquiry you want to modify, and press ENTER. The panels that were originally used to construct the existing inquiry are displayed. You can modify these panels to conform to your new requirements for the analysis inquiry and reexecute it. When you save the modified analysis inquiry, it replaces the original inquiry in your catalog.

1.4.3.3 Repeating an Existing Analysis Inquiry

This section describes how to use an existing analysis inquiry as a model for a new inquiry without changing the original inquiry. Figure 1-20 shows a typical display of the contents of a private catalog.

[illegible]

Figure 1-20. Typical Private Catalog - Repeat Example

To repeat a copy of one of these definitions in your private catalog, type R (Repeat) in the Cmd column on the line of the inquiry you wish to repeat and press ENTER. The Inquiry Repeat/Copy panel is shown in Figure 1-21.

```
/----- Inquiry Repeat/Copy -----Repeat Inquiry\
|Command ==>                               |
|                                           |
|Catalog: PRIVATE                          |
|                                           |
|    Inquiry Name ==> ____                 |
|           Title ==> Data Clustering      |
|           Author ==> CA MICS User         |
|                                           |
|Catalog Group Name ==> PERFORM             |
|           Title ==> Performance Management |
|                                           |
|                                           |
|Press ENTER to add the new inquiry to your private catalog.
|or
|Press END to cancel.
|                                           |
|                                           |
|                                           |
|                                           |
\-----/
```

Figure 1-21. Inquiry Repeat/Copy Panel

After the Inquiry Name prompt, enter a new inquiry name for this repeated analysis inquiry. The values for Title, Author, Catalog Group Name, and Title are copied from the inquiry being repeated. You can change them if you wish. When you are satisfied with the entries, press ENTER. The Performance Manager Catalog panel is redisplayed to allow you to select the repeated inquiry for modification, as discussed in Section 1.4.3.2. The status of the original inquiry analysis is not affected by this procedure.

1.4.3.4 Deleting an Analysis Inquiry

Figure 1-22 shows a typical display of the contents of a private catalog.

[illegible]

Figure 1-22. Typical Private Catalog - Delete Example

To delete one of these inquiries from your private catalog, type D (Delete) in the Cmd column of the analysis inquiry to be deleted and press ENTER. The Data Base Inquiries Confirm Delete panel is shown in Figure 1-23.

```

/-----Data Base Inquiries-----Confirm Delete\
| Command ==>|
|
| *****|
| * *|
| * Confirm Delete *|
| * *|
| * Press ENTER to DELETE or END to bypass *|
| * *|
| *****|
|
| Inquiry Name: WKLD1|
| Title: Data Clustering|
| Author: CA MICS User|
|
| Catalog: PRIVATE|
| Catalog group name: PERFORM|
| Title: Performance Management|
|
| \-----/

```

Figure 1-23. Data Base Inquiries Confirm Delete Panel

Press ENTER to confirm the delete request or type END to cancel the delete request.

1.4.3.5 Executing an Analysis Inquiry

To select an analysis inquiry for execution, type S (select) in the Cmd column of the inquiry on the main Performance Manager Catalog panel shown in Figure 1-24.

[illegible]

Figure 1-24. Selecting an Inquiry for Execution

The MICF Execution panel is shown in Figure 1-25.

```

/----- MICF Execution -----\
|Command ==>                               |
|                                           |
|Executing CA MICS Inquiry: WKLD1 -Testing fix for PE177130 - Invalid IF |
|                                           |
|Inquiry Execution Mode    ==> BATCH      (BATCH/FOREGROUND/DMS) |
|Override Execution Options ==> N        (Y/N) |
|SAS/DMS Options:  Resume ==> N        (Y/N) Output Replay ==> Y (Y/N) |
|                                           |
|Press ENTER to execute the inquiry, or press END to cancel execution. |
|                                           |
|----- Batch Job Submission Parameters -----|
|Edit Generated Jobstream ==> Y (Y/N) Hold Inq. Output for Replay ==> N (Y/N)|
|                                           |
|Sysout Class ==> *                               Sysout Form ==> ____ |
|Destination ==> _____ Sysout Writer ==> _____ |
|Sysout Parm. ==> _____ |
|                                           |
|Job Card Information: |
| ==> //jobname JOB (accounting information), |
| ==> //MSGLEVEL=(1,1),MSGCLASS=A,PRTY=5 |
| ==> /* |
| ==> /* |
\-----/

```

Figure 1-25. CA MICS Data Base Inquiry Execution Panel

The MICF Execution screen offers you the following execution options:

- o Selecting the Execution Mode
- o Overriding the Default Execution Options
- o Overriding the Default Job Submission Parameters

After you specify the options you desire and press ENTER, you are provided with a fifth option:

- o Selecting Execution-Time Data

Each of these options is discussed next.

- 1 - Selecting the Execution Mode
- 2 - Overriding the Default Execution Options
- 3 - Overriding the Default Job Submission Parameters
- 4 - Selecting Execution-Time Data

1.4.3.5.1 Selecting the Execution Mode

There are two execution modes: batch and foreground. In batch mode, the online facility creates and submits a batch job that executes the inquiry. In foreground mode, the online facility executes the inquiry interactively and lets you review the results immediately. To select an execution mode, type either BATCH (or B) or FOREGROUND (or F) after the Inquiry Execution Mode prompt.

If you select foreground, the inquiry executes interactively and gives you the ability to browse your output online, print it, and save it. The disadvantage of foreground execution is that you cannot use your terminal for anything else while an analysis inquiry is executing.

If you select batch, a batch job is generated and submitted to execute the inquiry. You have the option to hold outputs (that is, color graphics, printed reports, MICSLOG, and SAS logs) for later online review (and/or printing). Or you can let the batch job print reports and logs according to your sysout specifications. In either case, MICS can display color graphics on the hardcopy graphics device you specify.

Some Performance Manager Option inquiries executing in the foreground environment require that SAS be executed via the SASCP command rather than by using the CALL command. These inquiries create temporary PROC FORMATS in ddname LIBRARY during their execution. SASCP searches LIBRARY for temporary formats if the format is not found in ddname SASLIB. However, SAS's CALL mode builds a directory of existing formats within SASLIB files during its start-up processing and therefore will never find any temporary formats created afterwards during the inquiry processing.

The CA MICS Administrator specifies the SAS foreground interface mode along with other SAS Global parameters through the CA MICS Administrator Facility options. Refer to Section 4.4.2 of the CA MICS Planning, Installation, Operation, and Maintenance Guide (PIOM) for more information.

Figure 1-26 shows a comparison of the features available in the batch and foreground execution modes. The ability to interactively replay your batch execution output is only available if you set "Hold inq. output for Replay" to Y. Otherwise your batch output goes directly to the specified sysout class.

Feature	Batch	Foreground

Ability to catalog output	Yes	Yes
Ability to replay color graphics	Yes	Yes
Ability to print reports	Yes	Yes
Ability to print color graphics	Yes	Yes
Terminal available for use during inquiry execution	Yes	No

Figure 1-26. Comparison of Batch and Foreground Execution

1.4.3.5.2 Overriding the Default Execution Options

You may not need to override the default execution options very often. But you should familiarize yourself with the available options by entering Y after the Override Execution Options prompt. The Execution Options panel (Figure 1-27) lets you specify the color graphics format set, the color graphics device, and the SAS execution parameters for the current inquiry.

```

/----- Execution Options -----
|Command ==>
|
|Executing CA MICS Inquiry: WKLD1 - Data Clustering
|Inquiry Execution Mode: BATCH
|
|Color Graphics Format Set      ==> TERMINAL
|Display Graphics During Execution ==> N (Y/N)
|
|Color Graphics Device      ==> IBM32793  Override Device Parms.    ==> N (Y/N)
|
|List the SAS Source Code    ==> Y (Y/N)  Abend For Errors      ==> N (Y/N)
|List Expanded Macros        ==> N (Y/N)  Max. No. Error Msgs. ==> ____
|Print Notes in Listing      ==> Y (Y/N)  Minimum Free Memory   ==> ____ (kbytes)
|Overprint Errors            ==> N (Y/N)  Printer Char. Set     ==> 96 (48/60/96)
|
|Page Dimensions: Columns   ==> 80        (64 - 256)
|                      Lines ==> 23        (20 - 500)
|Exec. Limits: CPU Time     ==> _____ (seconds)
|                      Lines ==> _____
|                      Pages  ==> _____
|SAS Invocation Options     ==> _____
\-----

```

Figure 1-27. Execution Options Panel

Most of the parameters on this panel are self-explanatory; however, the Execution Options tutorial panels provide more detail if you should need it. You can reach these panels from the Execution Options panel by typing HELP at the Command prompt or (usually) by pressing your PF1 key.

There are two sets of execution parameters: one set is for batch job submissions and the other is for inquiries run in the foreground. These options are originally set by your CA MICS System Administrator, although you can personalize them by using Option 0, MICF Options, in the MICF facility. The online facility uses the set of options that corresponds to your choice of execution mode: batch or foreground.

1.4.3.5.3 Overriding the Default Job Submission Parameters

When you instruct the online facility to execute an inquiry as a batch job, it automatically uses your default job control specifications. These include your job statement and the SYSOUT specifications for the printed report and the SAS log.

You can override the default specifications for the current inquiry execution and modify batch job submission parameters and job card information from the MICF Execution panel (see Figure 1-25). If you want to edit the generated job stream and inquiry program before you submit the job, type Y at the Edit Generated Jobstream prompt.

If you specify Y (yes) following the "Hold Inq. Output for Replay" prompt, MICF will save inquiry outputs (i.e., color graphics, printed reports, MICSLOG, SAS log) in your private batch staging inquiry output catalog. You can use the MICF Reports and Graphics option to review inquiry outputs after the batch job completes execution. Specify N (no) for the hold batch outputs parameter if you want standard batch inquiry execution with hardcopy outputs.

You can replay color graphics and/or browse printed reports, MICSLOG, and SAS log from your batch staging catalog. You can move inquiry outputs from your batch staging catalog to your private inquiry output catalog. You can also delete batch staging catalog entries.

1.4.3.5.4 Selecting Execution-Time Data

The online facility analysis inquiries allow you to specify at execution time information such as the range of CA MICS file cycles or the system identifier (SYSID) to be reported on. When you press ENTER to select an analysis inquiry for execution, the online facility presents you with the Execution-Time Parameters panel.

Figure 1-28 illustrates an Execution-Time Parameters panel after the Data Base ID and CA MICS File Cycles have been entered. The content of this panel changes depending on the type of analysis you are executing, but the format is the same as that in Figure 1-28. To execute the analysis inquiry, type END after the Command prompt and press ENTER, or (usually) press your PF3 key.

```

/----- Execution-Time Parameters -----\
|Command ==>                               Scroll ==> CSR |
|Press END to execute the inquiry. Enter CANCEL to cancel inquiry execution. |
|Executing CA MICS Inquiry: WKLD1 - Data Clustering |
| |
|-----|
| Enter one DBID or a list of DBIDs for the data Base(s) you want to process. |
|A - _____|
| |
| Enter the range of File Cycles (e.g., 01 - 04) to be processed. |
|01 - 02|
| |
|***** BOTTOM OF DATA *****|
| |
| |
| |
| |
| |
| |
| |
|-----\

```

Figure 1-28. Execution-Time Parameters Panel

1.4.3.6 Specifying Common Extended Options

Supplied with the Reporting Analysis queries and Analysis tools queries are panels that you can use to specify additional options that are normally defaulted.

Typing a Y at the "Specify extended options" prompt on any of the query definition panels displays an Extended Options Menu. The common extended options are shown in Figure 1-29.

Note: These are the most common extended options; each query may not have all of these options. Extended options unique to a particular query are documented in the query's Component Operation section of this guide.

```
/----- Extended Options Menu -----\  
|Command ==>  
|  
|Inquiry Step:  Generic Performance Manager Query  
|  
|  
|Specify global data selection           ==> N (Y/N/R)  
|Specify alternate input file           ==> N (Y/N/R)  
|Save selected input observations        ==> N (Y/N/R)  
|Define alternative report elements      ==> N (Y/N/R)  
|Specify non-CA MICS data element derivation exit ==> N (Y/N/R)  
|Override color graphics format parameters ==> N (Y/N/R)  
|          Color graphics format set    ==> PENPLOT  
|  
|  
\-----/
```

Figure 1-29. Extended Options Menu

You can select any of the options available on this panel by typing a Y (YES) at its prompt. Typing an R (RESET) resets the option and nullifies any specifications you made for that option.

The following sections provide detailed instructions on how to use the different extended options.

- 1 - Global Data Selection
- 2 - Alternate Input File
- 3 - Save Selected Input Observations
- 4 - Define Alternative Report Elements
- 5 - Data Element Derivation
- 6 - Override Color Graphics Format Parameters

1.4.3.6.1 Global Data Selection

If you enter Y after Specify global data selection on the Extended Options Menu, the panel shown in Figure 1-30 appears.

```

/----- SAS Statements -----\
|Command ==>                      Scroll ==> CSR|
|
|Inquiry Step: Data Clustering
|
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy
|
|Cmd  SAS Statements for: Specifying global data selection
| - -----
|_ - %MACRO SELEOF;
|
|-----
|
|-----
|
|-----
|
|-----
|
|-----
|
|-----
|
|-----
|
|-----
|
|_ - %MEND SELEOF;
|
|***** BOTTOM OF DATA *****|
|
|-----\

```

Figure 1-30. Global Data Selection Panel

Several queries provide a user exit that you can use to control data selection from the CA MICS database. To invoke this exit, code SAS statements in the SELEOF macro. This macro is bracketed by the %MACRO and %MEND statements, and field-protected to prevent modification.

The default for the SELEOF macro is null. This macro is invoked when the workload observations are read from the specified CA MICS file. You can use it to delete observations from the data as it is selected. For example, in a job class study, you might want to exclude production job classes from the analysis. If production jobs are designated by job class P, then code the following:

```
%MACRO SELEOF;  
  IF JOBCLASS='P' THEN GO TO T_EOF;  
%MEND SELEOF;
```

Note that you do not have to code the %MACRO and %MEND statements, as they are already provided for you.

The IF statement should specify branching to T_EOF rather than acting as a subsetting IF or specifying DELETE. A subsetting IF or a DELETE statement may cause the program to function incorrectly.

You must code each specified SAS statement in accordance with the syntax of the SAS Macro Language, as described in the SAS manuals.

```
/----- Alternate Input File -----\
|Command ==>
|
|Inquiry Step: Data Clustering
|
|SAS file name(s) ==> _____
|
|
|SAS Data set allocation instructions:
|   Data set name    ==> _____
|   Disposition      ==> SHR  (SHR/OLD)
|
|
|
|
|
|
|
```

Several queries have the ability to read data from either CA MICS files or external SAS files. The Alternate Input File panel allows you to specify that the input data for the query is a SAS file or in some cases up to five SAS files. The option also allows you to specify the input SAS data set name. It is possible to specify an archive tape file, but only when the query does not require multiple input files. Note that the SAS data must already be created.

1.4.3.6.3 Save Selected Input Observations

If you enter Y after Save selected input observations on the Extended Options Menu, the panel shown in Figure 1-32 appears.

[illegible]

Figure 1-32. Save Selected Input Observations Panel

Several queries have the ability to save the observations selected by your analysis inquiry to an external SAS file for further processing by independent SAS programs, or for future processing by the analysis query. The Save Selected Input Observations panel allows you to specify that the selected observations are to be written to an external SAS file, and the SAS data set name and the name of the SAS file to which they are to be written. Note that the SAS data set must already be created.

1.4.3.6.4 Define Alternative Report Elements

If you enter Y after Define alternative report elements on the Extended Options Menu, the panel shown in Figure 1-33 appears.

[illegible]

Figure 1-33. Define Alternative Report Elements Panel

The Define Alternative Report Elements panel allows you to specify the data elements to be included on the report. To request a list of the elements in the input file, enter a question mark (?) in the Element field.

1.4.3.6.5 Data Element Derivation

If you require user-defined data elements, use the Data Element Derivation Macro Definition panel shown in Figure 1-34 to specify their derivation.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> CSR
|
|Inquiry Step:  Data Clustering
|
|Line Cnds: I Insert  D Delete  R Repeat  M Move  C Copy
|
|Cmd  SAS Statements for:  Specifying data element derivation
| -  -----
|_ - %MACRO  DRVMAC;
|
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
|_ - %MEND  DRVMAC;
|
|***** BOTTOM OF DATA *****
|
\-----

```

Figure 1-34. Data Element Derivation Macro Definition Panel

Data Clustering provides a user-exit for defining the procedures for computing data elements that are not part of the CA MICS database. This exit is invoked by coding SAS statements in the DRVMAC macro. This macro is bracketed by the %MACRO and %MEND statements, and field protected to prevent modification. The default for the DRVMAC macro is null.

You must code each data element derivation SAS statement in accordance with the syntax of the SAS Macro Language, as described in the SAS manuals.

1.4.3.6.6 Override Color Graphics Format Parameters

If you enter Y after Override color graphics format parameters on the Extended Options Menu, the Color Graphics Format Parameters panel, shown in Figure 1-35 appears.

You can use this panel to override parameter defaults that control color graphics format and appearance. These default values are specified through the MICF Color Graphics Format Parameters process. (SAS/GRAPH software is a prerequisite for color graphics.) You can customize these defaults to match your most common reporting requirements. In this way, you need to override color graphics format defaults only for exceptional reporting requirements.

```

/----- Color Graphics Format Parameters -----\
| Command ==>                                     Format set: PENPLOT |
| Inquiry step: IPS - Multiple Domain Analysis Plots |
| Title print font      ==> XSWISSBE   Title color      ==> BLUE |
| Draw border around plot ==> N   (Y/N) Background color ==> _____ |
| Axis color            ==> BLACK     Area inside axis color ==> _____ |
| Text color            ==> BLACK |
|
| Domain #1:  Color ==> RED           Sym ==> NONE           Line ==> SOLID |
| Domain #2:  Color ==> YELLOW        Sym ==> NONE           Line ==> SOLID |
| Domain #3:  Color ==> GREEN         Sym ==> NONE           Line ==> SOLID |
| Domain #4:  Color ==> CYAN          Sym ==> NONE           Line ==> SOLID |
| Domain #5:  Color ==> PINK          Sym ==> NONE           Line ==> SOLID |
| Domain #6:  Color ==> WHITE         Sym ==> NONE           Line ==> DOTTED |
| Domain #7:  Color ==> BLUE          Sym ==> NONE           Line ==> DOTTED |
| Domain #8:  Color ==> RED           Sym ==> NONE           Line ==> DOTTED |
| Domain #9:  Color ==> YELLOW        Sym ==> NONE           Line ==> DOTTED |
|
| Symbol (NONE,PLUS,STAR,TRIANGLE,DIAMOND,HASH,SQUARE,A-Z,0-9,or SAS/GRAPH value) |
| Color  (May be any color available on your graphics device.) |
| Line   (SOLID,DASHED,DOTTED,1-32) |
\-----/

```

Figure 1-35. Color Graphics Format Parameters Panel

The options for the Color Graphics Format Parameters panel are:

Title print font:	The character format (font) for the title lines. Character font names are one to nine characters in length. See the SAS/GRAPH Reference Guide for a list of valid character fonts.
Title color:	The color of the title lines.
Draw border around plot:	The specification to frame the graph with a border (that is, draw a line around the outside of the display).
Background color:	The specification to generate the graph on a colored background (that is, draw a line around the outside of the display).
Axis color:	The color of the x-axis and y-axis.
Area inside axis color:	The specification to draw the plot lines on a colored background (that is, only the area inside the axis is filled).
Text color:	The color of all text words except the titles and labels.
Symbol:	The symbol for plotting the data values. See the SAS/GRAPH Reference Guide for the available symbols.
Color:	The color for plotting the data values. See the SAS/GRAPH Reference Guide for the available colors.
Line:	The line type for plotting the data values. See the SAS/GRAPH Reference Guide for the available line types.

1.4.4 Replaying Cataloged Output

The online facility uses the MICF facilities for replaying, printing, and deleting cataloged output.

Chapter 2: PERFORMANCE REPORTING ANALYSIS

The CA MICS Performance Manager, Performance Reporting Analysis section provides interactive access to reports, and summarizes the system performance using the latest measurement information. This process is used to extract selected information from the CA MICS database and then apply analysis techniques and methodologies in preparing reports, charts, and graphs that facilitate the performance management process.

CA MICS Performance Manager operates under the CA MICS Information Center Facility (MICF) environment. MICF provides execution facilities to define, modify, or update performance analysis inquiries in the MICF inquiry catalog.

This section contains the following topics:

- [2.1 ILM Analysis](#) (see page 68)
- [2.2 Real Storage Management Analysis](#) (see page 89)
- [2.3 Auxiliary Storage Management Analysis](#) (see page 196)
- [2.4 I/O Configuration Analysis](#) (see page 251)
- [2.5 Virtual Storage Management Analysis](#) (see page 449)
- [2.6 PR/SM LPAR Performance Analysis](#) (see page 524)
- [2.7 Graphic Analysis](#) (see page 556)
- [2.8 Sysplex Analysis](#) (see page 613)
- [2.9 WLM Goal Mode Analysis](#) (see page 664)
- [2.10 Product Measured Usage Analysis](#) (see page 700)

2.1 ILM Analysis

This section introduces IBM License Manager (ILM) as a background for discussing the analysis of MSU (millions of service units) consumption rate, defined capacity, and soft-capping.

Section 2.1.4, ILM Concepts, defines the terms and explains the basic dynamics of the elements used in ILM analysis. Depending on your familiarity with these terms and concepts, you may find it useful to read this section before proceeding with Sections 2.1.2 and 2.1.3.

The following sections describe the operation of the MSU Consumption Rate Analysis inquiries and how to use the reports and graphs produced:

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output
- 4 - ILM Concepts
- 5 - Requirements

2.1.1 Functional Description

ILM analysis inquiries let you analyze the following items:

- o The MSU consumption rates of z/OS workloads
- o WLM's soft capping activity
- o Defined Group Capacity MSU consumption
- o The estimated MSU consumption of non-z/OS workloads

The graph produced when analyzing z/OS workloads plots MSU consumption rate, MSU four-hour rolling average, and the defined capacity for each PR/SM controlled LP present in the data. When defined group capacity analysis is specified two additional reports are generated, separately charting the summarized MSU consumption rate and four-hour rolling average of each capacity group.

The report produced when analyzing z/OS workloads contains the same metrics displayed in the graph with the addition of a soft-capping percentage for the LP and a threshold indicator. This field will indicate if the soft-capping percentage exceeded the user-specified soft-capping threshold.

You also have the option to activate threshold alerting. This produces a ServicePlus Service Desk alert when the soft-capping threshold has been exceeded.

For more information on using the ServicePlus Service Desk alert facility, see section 4.9 of the CA MICS Resource Management Planning, Installation, Operation, and Maintenance Guide.

MSU consumption rate analysis of non-z/OS workloads at the DETAIL timespan calculates the MSU consumption rate and the MSU four-hour rolling average of the non-z/OS workload as if the same workload were running under z/OS. At higher timespans, the ability to calculate the MSU four-hour rolling average is lost; therefore, these elements are dropped from the report.

The ILM analysis queries use data from the CPU Activity File (HARCPU). The range of data that can be reported from a particular CA MICS database depends on the timespans that were defined for this file and the number of cycles defined in each timespan. If you are not familiar with the generation

options for the database unit containing the MVS HAR Information Area, contact your CA MICS administrator for assistance.

2.1.2 Usage Guidelines

ILM allows PR/SM partitions to be configured with a defined capacity, expressed in MSUs. IBM's Workload Manager monitors the actual LP MSU utilization using a rolling four-hour average. If the MSU four-hour average consumption exceeds the defined capacity, WLM instructs PR/SM to "soft cap" the LP, restricting processor consumption until the four-hour average drops below the defined capacity. Thus, with ILM, an LP can deal with short-term spikes in processor utilization far above the defined capacity, as long as the four-hour rolling average remains below the defined capacity.

Group Capacity, a new function introduced with z/OS Version 1 Release 8, allows customers to apply the service units defined capacity limit to not only a single logical partition (LP) but also to a group of LPs running on the same server.

This section discusses the application of the concepts discussed in Section 2.1.4, ILM Concepts. The following queries will be covered in more depth:

- 1 - MSU Consumption Rate and Soft Capping Analysis
- 2 - Estimating MSU Consumption of non-z/OS Workloads

2.1.2.1 MSU Consumption Rate and Soft Capping Analysis

The MSU Consumption Rate and Soft Capping plot and report analyzes HARCPU data collected from z/OS workloads. Using detail-level data, this query tracks MSU consumption rate, four-hour rolling averages, and the percentage of time an LP has been soft-capped. These metrics are plotted along with the total number of SUs available to the partition. When defined capacity limits are implemented, the SUs available to the partitions reflects the defined capacity limit. The tabular report lists these metrics along with a soft capping threshold indicator.

The soft capping percentage threshold can be specified on the extended options panel. The query uses a default of 05%, unless otherwise specified. The threshold indicator is populated with an eye catcher of "<- EXCEEDED 05%" when the user-specified soft capping percentage is exceeded. The MICSLOG message PER03900I indicates at a glance if any LPARs exceeded the threshold.

The query also lets you generate a ServicePlus Service Desk alert when the soft capping threshold has been exceeded. This can be done by activating threshold alerting on the extended option panel. This feature works in conjunction with the soft capping threshold extended option. The MICSLOG message PER03900I indicates at a glance if any LPARs exceeded the threshold, and the MICSLOG BAS63099I message indicates if the Service Desk alert was generated and the request number for reference.

When analyzing the report generated by this query, the following are some guidelines to follow:

- o When the defined capacity is exceeded by the MSU four-hour rolling average.

If the MSU four-hour rolling average exceeds the defined capacity for long periods (high WLM capping percentage reported), you may consider performing a workload analysis on the four hours leading up to this event, to help you understand the workload associated with the spiking in MSU consumption rate. You may also consider looking at the performance impact that the soft capping had on your performance groups.

- o When the MSU four-hour rolling average and the WLM capping percentage values are always zero while the MSU consumption rate is non zero.

When this occurs, it signifies that you are processing z/OS data but ILM's soft capping has not been activated. The LP in question has implemented either a hard cap or an uncapped LP.

- o Setting your defined capacity

Ideally, your peak MSU consumption rate must always be as close as possible to the defined capacity, without exceeding it. A general rule to follow when setting your defined capacity is to set it near the MSU consumption rate's 95th percentile.

Remember that capping affects all workloads running in the partition and can result in elongated response times and generally unacceptable performance.

The following are some additional guidelines to follow once you have set the defined capacity:

- o If the peak MSU consumption rate is constantly lower than the defined limit, you may consider decreasing the LPAR capacity until WLM begins capping the partition. (A capping percentage of 1% to 5% ensures the work is done optimally at the lowest cost.)
- o If the peak value exceeds the defined capacity for long periods (high WLM capping percentage reported), you may consider increasing the LPAR capacity limit until the WLM capping percentage falls to within an acceptable range.
- o If the peak value occasionally exceeds the defined capacity (low WLM capping percentage), you may consider time-shifting the workload.

2.1.2.2 Estimating MSU Consumption for non-z/OS Workloads

The Estimated MSU Consumption Rate Report analyzes historical HARCPU data collected from non-z/OS workloads, calculates the max MSU consumption rate, the average MSU consumption rate, and the MSU four-hour rolling average. These estimations are accomplished by extracting the processor model and the number of processors from the historical data, and pulling the associated MSU conversion factor from a provided table. For further discussion on how this MSU conversion factor is used, see section 2.1.3.2.

When analyzing the report generated by this query here are some guidelines to follow:

- o Setting your defined capacity

Ideally, your peak MSU consumption rate should always be as close as possible to the defined capacity, without exceeding it. A general rule to follow when setting your defined capacity is set it near the MSU consumption rate's 95th percentile.

Remember capping impacts all workloads running in the partition and can result in elongated response times and generally unacceptable performance. The following are some additional guidelines to follow once you have set the defined capacity.

- o If the peak MSU consumption rate is constantly lower than the defined limit, you may consider decreasing the LPAR capacity until WLM begins capping the partition (a capping percentage of 1% to 5% ensures the work is done optimally at the lowest cost).
- o If the peak value exceeds the defined capacity for long periods (high WLM capping percentage reported), you may consider increasing the LPAR capacity limit until the WLM capping percentage falls to within an acceptable range.
- o If the peak value occasionally exceeds the defined capacity (low WLM capping percentage), you may consider time shifting workloads.
- o Estimated MSU Four-Hour Rolling Average values are zero

Due to the nature of a rolling average, you need to process a certain amount of data before you can begin calculating

the rolling average. In the case of the Estimating MSU Consumption Rate report the MSU four-hour rolling average will be zero until you have processed four hours worth of data.

This becomes important to remember if you decide to report on an individual zone. If the zone consist of five hours, the report will display zero's for the MSU Four-Hour Rolling Average for the first four hours and only give you values for the last hour.

o Warning message PER00121W is issued in the MICSLOG

Warning message PER00121W is issued when the historical input data contains combination of processor model and number of processors not defined in the MSU Conversion Table. If this message is issued update, the MSU Conversion Table through the extended options panel. To update this table you will be required to have the processor model, the number of processors, and MSU rate associated to it.

2.1.3 Standard Output

COMMON REPORT ELEMENTS

Each report has several lines of title information that give the name of the report, the installation name, the run date, the date(s) being reported, and the start and end times of the period covered by the report. The following elements are in the standard title part of each report:

Installation Name: On some reports, the installation name that is defined at CA MICS system installation time. This value is centered in the first line of the report. For more information about this value, consult your CA MICS System Administrator.

Component Name: The name of the performance component, centered in the first or second line of each page.

Report Name: The name of the report, centered in the third line of the heading. It cannot be modified.

SYSID or System: The CA MICS logical SYSID being reported. Data from multiple systems cannot be combined using standard control parameters.

ZONE: If data is being reported by zone, the zone number is shown.

RUNDATE: The date the report was run.

REPORT OPTIONS

The CA MICS Performance Manager reports are initially delivered with defaults. These defaults will change based on the settings of the last execution of the query. For example, the Estimated MSU Consumption report is delivered with a hard-coded default value of DETAIL for the timespan. If the user changes the timespan so that report will run at the WEEKS timespan, then the next execution of this report will default to the WEEKS timespan.

The Report supports standard time selection options, as well as SYSID selection and external data selection. These options are controlled by the control options on the execution panel.

All graphs and reports can be produced at multiple summarization levels. Dates and titles will be adjusted

appropriately based on the timespan and summarization level chosen for the execution. After data has been collected for a time period, this summarization capability can be used to "drill down" to a specific time frame.

The following analysis queries are discussed in more detail:

- 1 - MSU Consumption Rate and Soft Capping Analysis
- 2 - Estimating MSU Consumption of non-z/OS Workloads

2.1.3.1 MSU Consumption Rate and Soft Capping Analysis

This query analyzes HARCPU data collected from z/OS workloads. Using detail-level data, this query tracks MSU consumption rates, four-hour rolling averages, and the percentage of time an LP has been soft-capped. When defined group capacity analysis is specified, two additional reports are generated, separately charting the summarized MSU consumption rate and four-hour rolling average of each capacity group. At higher levels of summarization, this query graphs the maximum MSU consumption rate and the maximum four-hour rolling average for the reporting interval. Also included are the service units available to z/OS image (CPUSUAVL), expressed in MSUs per hour.

Each report can be selectively produced using the Report Option on the query composition panel. The ability to produce a CA MICS Q&R CSV file is also available by selecting the Generate Q&R CSV File option. When selected, a CSV file will be produced containing all elements used in each report. The default view in Q&R will be MSU Consumption over time.

REPORT ELEMENTS

Report Label	File	Element
-----	----	-----
Max MSU Four-Hour Rolling Average	HARCPU	CPUMXLAC
Max MSU Consumption Rate	HARCPU	CPUMXMSU
WLM Capping Percentage	HARCPU	CPUPCNSW
SUs Available To z/OS Image	HARCPU	CPUSUAVL
Capacity Group Name	HARCPU	PRSMCGNM
Capacity Group Limit	HARCPU	PRSMCGLI

MSU CONSUMPTION RATE AND MSU FOUR-HOUR ROLLING AVG. REPORTS

To help you better understand the standard output produced by this query, definitions and calculations involved in producing the reported variables follow:

CEC The number of service units (expressed in
CAPACITY MSUs) available to the partition, that is the
IN MSUs: LP CPU capacity.

If an LP has no defined capacity limit, CPUSUAVL is based on the number of logical processors defined for the partition if the LP is uncapped; or, on the LP's weight, if it is capped via the Hardware Management Console.

If an LP has a defined capacity, CPUSUAVL contains this defined capacity limit, and is equivalent to the LPCDEFLM data element of the HARLPC (PR/SM LP Config/Activity) file.

MAX MSU CONSUMPTION RATE: Contains the total number of MSUs consumed by the partition, during the measurement interval.

$$\text{CPUTOMSU} = \frac{\text{CPUTODTM} * \text{CPUSUSEC} * 3600}{\text{DURATION} * 1000000}$$

CPUTODTM is the total dispatch time of all processors assigned to the partition.

CPUSUSEC is the physical CPU adjustment factor used to convert CPU times to SUs.

DURATION is the length of the RMF recording interval.

MAX MSU FOUR-HOUR ROLLING AVERAGE: Contains the maximum value of the four-hour rolling average number of service units consumed (expressed in MSUs) by the partition, during the CA MICS reporting interval.

WLM CAPPING PERCENTAGE: Contains the percentage of the time when the partition has been capped by WLM because the LP used more service units than its defined capacity limit, based on a MSU four-hour rolling average consumption.

CAPACITY GROUP LIMIT: The maximum number of processor licensing units for the group of logical partitions identified by having the same group name, of which this partition is a member. The capacity group limit is expressed in MSUs.




Figure 2-1. Sample MSU Consumption Rate and Soft Capping Report

MSU CONSUMPTION RATE AND SOFT CAPPING REPORT					
----- Sysplex Name=AEPLEX0A System Identifier=AE91 Logical Partition Name=AE91 Date=mm-dd-yy -----					
End Time	Max MSU Consumption Rate (CPUMXMSU)	Max MSU Four-Hour Rolling Average (CPUMXMSU)	SUs Available To MVS Image (CPUSUAVL)	WLM Capping Percentage (CPUPCNSW)	Soft Capping Threshold Indicator
10:00:00	35	20	90	0.00 %	
10:10:00	42	23	90	0.00 %	
10:20:00	55	27	90	0.00 %	
10:30:00	75	30	90	0.00 %	
10:40:00	65	33	90	0.00 %	
10:50:00	85	36	90	0.00 %	
11:00:00	104	40	90	0.00 %	
11:10:00	106	46	90	0.00 %	
11:20:00	110	55	90	0.00 %	
11:30:00	115	60	90	0.00 %	
11:40:00	110	65	90	0.00 %	
11:50:00	95	68	90	0.00 %	
12:00:00	115	69	90	0.00 %	
12:10:00	121	71	90	0.00 %	
12:20:00	115	73	90	2.33 %	
12:30:00	120	78	90	3.50 %	
12:40:00	130	84	90	3.70 %	
12:50:00	125	88	90	4.66 %	
13:00:00	91	90	90	7.32 %	<- EXCEEDED 07%
13:10:00	93	92	90	12.66 %	<- EXCEEDED 07%
13:20:00	91	90	90	15.33 %	<- EXCEEDED 07%
13:30:00	91	91	90	11.23 %	<- EXCEEDED 07%
13:40:00	100	90	90	8.21 %	<- EXCEEDED 07%
13:50:00	95	85	90	6.50 %	
14:00:00	99	85	90	4.33 %	
14:10:00	55	82	90	2.17 %	
14:20:00	65	79	90	0.00 %	
14:30:00	75	65	90	0.00 %	
14:40:00	55	63	90	0.00 %	
14:50:00	51	55	90	0.00 %	
CA MICS Performance Manager MSUXX 05FEByy Based on DETAIL LEVEL INFORMATION					

Figure 2-2. MSU Consumption Rate and Soft Capping Report




Figure 2-3. Sample MSU Consumption Rate by Capacity Group Report




Figure 2-4. Sample MSU Four-Hour Rolling Average by Capacity Group

2.1.3.2 Estimated MSU Consumption of non-z/OS Workloads

The Estimated MSU Consumption Rate Report analyzes HARCPU data collected from non-z/OS workloads and calculates the max MSU consumption rate, average MSU consumption rate, and the MSU four-hour rolling average.

When analyzing HARCPU data at the detail level timespan, this query calculates and reports the estimated maximum MSU consumption rate, estimated average MSU consumption rate, estimated MSU four-hour rolling average, and the CEC Capacity In MSUs.

Note: At higher timespans you lose the ability to calculate a MSU four-hour rolling average. Therefore, this element will be dropped from the report.

DATA ELEMENTS

Report Elements:

Report Label	File	Element
-----	----	-----
CEC Capacity In MSUs	calculated	MODMSU
Estimated Max. Consumption Rate	calculated	MSUMXPBS
Estimated Avg. Consumption Rate	calculated	MSUPCBSY
Estimated MSU Four-Hour Rolling Avg.	calculated	AVGMSU

Input Elements:

Element Label	File	Element
-----	----	-----
CPC Model Identifier	HARCPU	CPCMODID
Processor Model Family	HARCPU	CPUMOD
Max Pct Processor Busy	HARCPU	CPUMXPBS
Max Number of Configured Proc	HARCPU	CPUNOCPU
Pct Processor Busy	HARCPU	CPUPCBSY
Total Time Processor Busy	HARCPU	CPUTOBTM
Processor Version	HARCPU	CPUVER
PR/SM LPAR Central Complex	HARCPU	PRSMCEC
Logical Partition Number	HARCPU	PRSMLPNO
Number of Physical Processors	HARCPU	PRSMNOPP

ESTIMATED MSU CONSUMPTION RATE REPORT

To help you better understand the graphs and reports produced by this query, definitions and calculations for the graphed variables follow:

CEC
CAPACITY
IN MSUs: The number of service units(expressed in
millions of service units per hour, MSUs)
available to the partition, i.e. the LP CPU
capacity.

ESTIMATED
AVERAGE MSU
CONSUMPTION
RATE: Contains the four-hour rolling average number
of service units consumed (expressed in MSU's)
by the partition during the CA MICS reporting
interval.

$$\text{MSUPCBSY} = (\text{CPUNOCPU}/\text{PRSMNOPP}) * (\text{CPUPCBSY}/100) * \text{MODMSU}$$

where: CPUNOCPU is the total dispatch time of all
processors assigned to the partition

PRSMNOPP is the physical CPU adjustment factor
used to convert CPU times to SU's

CPUPCBSY Pct Processor Busy

MODMSU is the CEC Capacity In MSUs

ESTIMATED
MAX. MSU
CONSUMPTION
RATE: Contains the maximum value of the four-hour
rolling average number of service units consumed
(expressed in MSU's) by the partition, during
the CA MICS reporting interval.

$$\text{MSUMXPBS} = (\text{CPUNOCPU}/\text{PRSMNOPP}) * (\text{CPUMXPBS}/100) * \text{MODMSU}$$

where: CPUNOCPU The number of processors present in the
configuration (online or offline) at
the end of the measurement interval.
This represents only tightly-coupled
CPU's and has a value of 1 for a UP,
2 for MP/AP and dyadic, 4 for tetradic.

PRSMNOPP is the number of physical processors
assigned for use by PR/SM.

CPUMXPBS is the maximum percent busy for the
processors at any time during the
reporting interval.

MODMSU is the CEC Capacity In MSUs

ESTIMATED MSU CONSUMPTION RATE REPORT				06:27 Monday, February 4, yyyy
----- Sysplex Name=SYSPLEXE System Identifier=SYS2 Logical Partition Name=XAD1 Date=08-02-26 -----				
End Time	Estimated Average Consumption Rate	Estimated Maximum Consumption Rate	Estimated MSU Four-Hour Rolling Average	CEC Capacity In MSUs
0:15:00	21	22	0	76
0:29:00	20	20	0	76
0:44:00	17	17	0	76
0:59:00	16	16	0	76
1:15:00	24	24	0	76
1:29:00	26	27	0	76
1:45:00	26	27	0	76
1:59:00	30	31	0	76
2:15:00	30	31	0	76
2:29:00	28	28	0	76
2:44:00	28	29	0	76
2:59:00	28	29	0	76
3:15:00	35	35	0	76
3:29:00	25	26	0	76
3:44:00	26	27	0	76
4:00:00	28	29	26	76
4:15:00	18	19	25	76
4:29:00	17	18	25	76
4:45:00	23	23	26	76
4:59:00	26	26	26	76
5:14:00	30	31	27	76
5:29:00	23	23	26	76
5:44:00	18	18	26	76
6:00:00	18	19	25	76
6:15:00	25	26	25	76
6:29:00	26	27	25	76
6:44:00	20	20	24	76
6:59:00	17	18	23	76
7:15:00	17	17	22	76
7:30:00	16	16	22	76
7:44:00	19	20	21	76
7:59:00	17	17	21	76
8:15:00	18	19	21	76
8:30:00	15	16	21	76
8:44:00	14	14	20	76
8:59:00	13	14	19	76
CA MICS Performance Manager JKJJC 04FEB02 Based on DETAIL LEVEL INFORMATION				
MSU CONSUMPTION RATE ESTIMATION REPORT				06:27 Monday, February 4, yyyy

Figure 2-4. Sample Estimated MSU Consumption Report

2.1.4 ILM Concepts

z/OS has introduced the concept of (ILM) IBM License Manager. ILM allows PR/SM partitions (LPs) to be configured with a defined capacity, expressed in MSUs per hour. IBM's Workload Manager (WLM) monitors the actual LP MSU utilization using a rolling four-hour average. If the MSU four-hour rolling average utilization exceeds the defined capacity, WLM instructs PR/SM to "soft cap" the LP, effectively restricting processor consumption until the four-hour average drops below the defined capacity. Thus, with ILM, an LP can effectively deal with short-term spikes in processor utilization far above the defined capacity, as long as the four-hour rolling average remains below the defined capacity.

Some software licenses are based on the defined capacity of the machine. Sites using ILM should closely track their peak four-hour rolling average values as well as the WLM capping percentage to correctly size their LP's capacity limits and, hence, their ability to achieve savings on software license charges by tailoring their license certificates.

Correctly sizing your LPARs capacity limits (and your ability to achieve savings on software license charges by tailoring your license certificates) requires the tracking of your peak four-hour rolling average values as well as the WLM capping percentage. The CA MICS data elements that report these values are, respectively, CPUMXLAC and CPUCNSW.

If the MSU consumption rate is constantly lower than the defined limit, you may consider decreasing the LPAR capacity until WLM begins capping the partition (a capping percentage of 1% to 5% ensures the work is done optimally at the lowest cost).

Capping affects all workloads running in the partition and can result in elongated response times and generally unacceptable performance.

If the peak value exceeds the defined capacity for long periods (high WLM capping percentage reported), you may consider increasing the LPAR capacity limit until the WLM capping percentage falls to within an acceptable range.

If the peak value occasionally exceeds the defined capacity (low WLM capping percentage), you may consider time-shifting the workload.

Ideally, your peak value should always be as close as

possible to the defined capacity, without exceeding it, while your WLM capping percentage should always be near to zero.

2.1.5 Requirements

Review the applicable sections to determine the execution requirements for the individual ILM performance reports.

- 1 - Requirements for ILM Analysis of z/OS Workloads
- 2 - Requirements for Estimating MSU Consumption Rates of non-z/OS Workloads

2.1.5.1 Requirements for ILM Analysis of z/OS Workloads

The following are requirements that you must satisfy in order to execute this report.

- o Processing RMF data from a z/OS platform
- o CPU Activity File (HARCPU) activated, with the following elements populated:
 - CPUMXMSU - Maximum Number of Consumed MSUs
 - CPUMXLAC - Maximum Long-Term Average of CPU MSUs
 - CPUSUAVL - Service Units Available to MVS Image
 - CPUPCNSW - Partition WLM Capping Percentage

Optional Defined Capacity Group Analysis Elements:

- PRSMCGNM - Capacity Group Name
- PRSMCGLI - Capacity Group Limit
- CPUCGJTS - Capacity Group Join Time Stamp
- CPUCGJCT - Intvls with Capacity Group Changed

Only one level is required for the CA MICS RMF product files (that is, DETAIL, DAYS, WEEKS, or MONTHS), depending on the timespan you specify during inquiry definition.

The number of cycles that are maintained at your installation for the CA MICS files determines the time frame for which you can produce ILM Analysis reports.

2.1.5.2 Estimating MSU Consumption of non-z/OS Workloads

The following are requirements that you must satisfy in order to execute this report.

In order to execute this report, you must have:

- o RMF data version 3.5 and above
- o Processing data collected from one of the following machine types:

G3, G4, G5, G5 Turbo, G6, G6 Turbo, Z900 Models

NOTE: If processing data from a machine type not listed above, you will be required to update the MSU conversion table through the extended options panel.

- o CPU Activity File (HARCPU) activated, with the following elements populated:

CPCMODID - CPC Model Identifier
CPUMOD - Processor Model Family
CPUMXPBS - Max Pct Processor Busy
CPUNOCPU - Max Number of Configured Processors
CPUPCBSY - Pct Processor Busy
CPUTOBTM - Total Time Processor Busy
CPUVER - Processor Version
PRSMCEC - PR/SM LPAR Central Complex
PRSMLPNO - Logical Partition Number
PRSMNOPP - Number of Physical Processors

Only one level is required for the CA MICS RMF product files (i.e., DETAIL, DAYS, WEEKS, or MONTHS), depending on the timespan you specify during inquiry definition.

The number of cycles that are maintained at your installation for the CA MICS files determines the time frame for which you can produce ILM Analysis reports.

2.2 Real Storage Management Analysis

This section introduces virtual storage concepts as a background for discussing the management of real storage in MVS. MVS virtual storage composition is described, with the goal of defining those parts of the virtual address space that require real memory.

Real Storage Management Analysis inquiries let you track the use of real storage by MVS at any supported timespan and summarization pattern. Reports display system measurements on real storage and swapping. In addition to system level reporting, you can display resource distribution patterns for such resources as frame utilization, using the private inquiry.

Section 2.2.4, MVS Concepts, defines the terms and explains the basic dynamics of many of the elements used in real storage analysis. Depending on your familiarity with these terms and concepts, you may find it useful to read this section before proceeding with Sections 2.2.2 and 2.2.3.

The following sections describe the operation of the Real Storage Management Analysis inquiries and how to use the reports that are produced:

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output
- 4 - MVS Concepts
- 5 - Analytic Tutorial
- 6 - Case Study
- 7 - Component Operation
- 8 - Requirements

2.2.1 Functional Description

The Real Storage Management Analysis inquiries allow you to track the utilization of real storage, resource (virtual storage) requirements by workload, and system swapping patterns. Any supported CA MICS timespan and summarization pattern can be implemented.

The System Storage Usage Report uses data from the MVS Paging Activity File (SCPPAG). Swap reports use data from the CA MICS MVS Swapping Activity File (SCPSWP). Private reporting as shown in the Private Storage Usage Report examples base the input file requirements on specifications you make during inquiry definition. Private reporting can be performed on any one of the two RMF workload files shown in the table below. Each file listed is comprised of workload data representing a defined workload level: Service Class, Service Class Period, and so on. The selection of the input file during inquiry definition also determines the workload level you wish to analyze.

WLMSEC - Service Class Resource Consumption
WLM_EC - Report Class Resource Consumption

The range of data that can be reported from a particular CA MICS database depends on the timespans that were defined for these files and the number of cycles defined in each timespan. If you are not familiar with the generation options for the database unit containing the MVS SCP Information Area, ask your system administrator for help.

2.2.2 Usage Guidelines

With the advent of z9xx hardware and z/OS system support, large memory models in the multi-gigabyte range have become commonplace. Accordingly, real storage management has mainly been reduced to simple monitoring to avoid high paging rates. Use of previous methods, such as storage isolation to protect sensitive workloads like CICS regions, has been superseded by WLM Goal Mode, where the actual monitoring and adjustment functions have been included in the WLM design. Accordingly, some sections of the CA MICS Performance Management documentation have been removed.

2.2.3 Standard Output

The RSM storage analysis reports give you the capability to track the allocation of real storage at any supported timespan and summarization level. This data is compared with the paging and swapping workload to allow analysis of the relationship between real storage allocation and the operation of the page/swap subsystem.

Real storage requirements at the RMF workload level (control or report performance group, domain, or performance period) are reported to help you evaluate the impact of new applications or users, as well as the growth of existing applications. You can use the selection parameters described in Section 2.2.7, Component Operation, to limit reporting to specific systems, times, and workloads.

Described below are the report formats of the System Storage Usage Report, the Private Storage Usage Report, the Swap Analysis Reports, and the Working Set Analysis Reports and Plots. Common title and report elements are also defined.

- 1 - Common Report Elements
- 2 - Real Storage System Usage Report
- 3 - Real Storage Private Usage Report
- 4 - Real Storage Swap Reason Report
- 5 - Real Storage Swap Type Report

2.2.3.1 Common Report Elements

This section describes elements common to all RSM Analysis reports. Options that are unique to a specific report are documented with that report.

TITLE ELEMENTS

Each report has five lines of title information available. All reports give the name of the report, the name of the component, and the run date. In addition, the heading on the System Storage Usage and Private Storage Usage Reports details the data selected, the summary level of the report, and the time periods reported by column:

- o For the monthly report, the column title is month/year.
- o For the weekly report, the first day of the week is given.
- o For the daily report, the day is given as month/day/year.
- o For the hourly level, the hour is given as the column title and the day of the month is given in the report title.

The following elements are in the standard title portion of each report:

Installation Name:	On some reports, the data center name that is defined at CA MICS system installation time. This value is centered in the first line of the report. For more information about this value, ask your system administrator.
Component Name:	The name of the performance component, centered in the first or second line of each page.
Report Name:	The name of the report, centered in the third line of the heading. It cannot be modified.
SYSID or System:	The CA MICS logical SYSID being reported. Data from multiple systems cannot be combined using standard control parameters.
ZONE:	If data is being reported by zone, the zone number is shown.
RUNDATE:	The date the report was run.

ADDITIONAL ELEMENTS

Reporting Interval:	The first row of some reports shows the date, or date/time combination of the period being reported. If the date being reported on is from the WEEKS timespan, the date reflects the starting date of the week, as derived from the CA MICS variable STARTTS. In all other cases, the date reflects the ending date, derived from the CA MICS variable ENDTS. In monthly reports, only the year and month portions of the date are displayed.
% CHANGE:	The last column on some reports, this field is the percent change for each report item (row) between the base period (first data column) and the last period (last data column). If data for only one period is reported, then dashes (---) are displayed. If the base period is zero, then N/A is printed to indicate that a percent change could not be calculated.

2.2.3.2 Real Storage System Usage Report

The Real Storage System Usage Report tracks the allocation of real storage frames at a monthly, weekly, daily, or hourly level. Figure 2-5 shows an example of a monthly Real Storage System Usage Report and Figure 2-6 shows the hourly form of the report. In the hourly form, data for one day is displayed by hour using the DAYS timespan of the MVS Paging Activity File (SCPPAG).

REPORT OPTIONS

The Real Storage System Usage Report supports standard time selection options, as well as SYSID selection and external data selection. These options are controlled by the control options discussed in Section 2.2.7.

DATA ELEMENTS

The elements reported are divided into the following categories:

- o Real Storage Use
- o Real Storage Availability
- o Fixed Storage Use
- o Fixed Storage Availability
- o Auxiliary Storage Workload
- o Expanded Storage Workload

REAL STORAGE USE

Real storage use is shown in terms of frames allocated to the MVS nucleus, the private area, the Local System Queue Area, the Common Storage Area, and the System Queue Area.

NUCLEUS FRAMES:	The size of the nucleus area, including the nucleus load module, nucleus control block (UCBs), and (for pre-MVS/XA systems) the fixed LPA and BLDL areas. This value comes from PAGFIN in the SCPPAG File.
PRIVATE FRAMES:	The total frames, both pageable and fixed, that are assigned to the private areas of all MVS memories. This is PAGAVS from the CA MICS database.
LSQA FIXED FRAMES:	The number of frames assigned to the Local System Queue Area (LSQA) of all MVS memories.

Because LSQA is totally fixed storage, the number of LSQA frames is also the amount of virtual pages used. Therefore, you can use this value to track LSQA allocation problems. The number of LSQA frames is PAGAVLS from the CA MICS database.

CSA FRAMES: The total frames assigned to the Common Storage Area (CSA). This value is PAGAVP from the CA MICS database.

LPA FRAMES: The total frames assigned to the Pageable Link Pack Area. This value is PAGAVLP from the CA MICS database.

SQA FIXED FRAMES: The total frames assigned to the System Queue Area (SQA). Because SQA is totally fixed storage, the number of SQA frames is also the amount of virtual pages used. You can use this value to track SQA growth problems. The value is PAGAVQ from the CA MICS database.

REAL STORAGE AVAILABILITY

Real storage availability is shown in terms of the total available frames and the average number of unused frames.

FRAMES IN REAL STORAGE: The number of frames available to the system other than the nucleus frames (PAGAVTFC). Variations in this number can mean that memory has been varied offline or become unavailable.

AVAILABLE FRAME COUNT: The average number of unused frames (the length of the ASM free queue). RMF only samples this number at the end of each measurement interval. It is included on this report as an indication of how real storage is being allocated. You should use the SCPSRCnn Files for analysis of the available frame queue. This value is PAGAVF

FIXED STORAGE USE

Real storage problems often arise from a growth in the amount of fixed storage, which effectively lowers the number of frames available for paging. To aid in tracking, all components of fixed storage are reported by this section of the report, even though several of the elements are reported in the first section. The pool of dynamic pages is further

reduced by storage isolation. This reduction is not reflected in this report.

NUCLEUS FRAMES:	Because the nucleus is fixed, it is included again in the fixed area of the report.
PRIVATE FIXED FRAMES:	The total fixed frames assigned to the private areas of all MVS memory frames. This is PAGAVR from the CA MICS database.
CSA FIXED FRAMES:	The total fixed frames assigned to the Common Storage Area (CSA). This value is PAGAVC from the CA MICS database.
LSQA FIXED FRAMES:	Because LSQA is fixed, the number of pages is also reported in the fixed area of the report.
SQA FIXED FRAMES:	Because SQA is fixed, the number of pages is also reported in the fixed area of the report.
LPA FIXED FRAMES:	The number of fixed pages assigned to the Pageable Link Pack Area (PAGAVLF).
FIXED BELOW 16 MB (XA):	The average number of real storage frames addressed between 0 and 16 MB, and marked non-pageable. This field is not valid for pre-MVS/XA systems.

FIXED STORAGE AVAILABILITY

Fixed storage availability is reported in terms of the total number of fixed frames and the number of frames that are either free or not fixed.

TOTAL FIXED FRAMES:	The total number of fixed frames over the report interval. This number is not the total of the separate pieces, as the value for each part of fixed storage is an average. This value is PAGAVX from the SCPPAG File.
TOTAL PAGEABLE FRAMES:	The number of frames that are either free (on the available frame queue) or are not fixed. This number is the base frame count that ASM has to work with. However, the minimum number of storage fenced pages must also be subtracted out. This value is PAGAVT from the CA MICS database. Because of changes to RMF in MVS/SP 1.3 systems and beyond, the number of pageable

frames is no longer recorded, resulting in the value for this field being the same as the FRAMES IN REAL STORAGE value.

AUXILIARY STORAGE WORKLOAD

The workload of the page/swap subsystem is reported so that you can view the relationship between the amount of pageable storage and page/swap rates at the various report levels.

DEMAND PAGE RATE:	The number of non-VIO, non-swap page-ins plus the non-VIO reclaims. This is the page fault rate (or demand page rate). This value is PAGPSDPG from the CA MICS database.
LPA PAGES PER SECOND:	The number of page-ins per second for the Link Pack Area. The LPA is not modified, so there are no page-outs for it.
CSA PAGES PER SECOND:	The number of page-ins plus page-outs per second for the Common Storage Area.
MAX DEMAND PAGE RATE:	The maximum observed demand paging rate over the time period (PAGMSDPG).
BLOCK IN RATE:	The rate of page blocks moving in from auxiliary storage (PAGPSBLK).
PAGES PER BLOCK IN:	The average number of pages per block moved. The value is: $(PAGPBLP + PAGBLK) / PAGBLK$
SWAP RATE:	The average number of physical swaps over the time period (PAGPSSWP).
MAX SWAP RATE:	The maximum swap rate over the time period (PAGMSSWP).
PAGES PER SWAP IN:	The average number of page frames involved per swap-in operation (PAGAVPSI).
TOTAL PAGES PER SECOND:	The total pages moved per second by the Auxiliary Storage Manager. This includes all physical movement of pages (swapping, demand, both in and out, and block paging) but no reclaims. This number is essentially the I/O overhead of the ASM. The value is:

(PAGPIN+PAGPOUT+PAGSIN+PAGSOT+PAGVIN+
PAGVOT+PAGBLP)/DURATION

PAGES The number of pages moved from one frame of real
MOVED/SEC: storage to another. For example, pages moved
 above or below the 16 MB line. Movements of
 pages to expanded or auxiliary storage are not
 included in this count.

VIRTUAL:REAL The ratio of allocated local page data set slots
RATIO: to the number of private area fixed and
 pageable frames.

EXPANDED STORAGE WORKLOAD

If your system is at or above RMF version 3.3.0 and the system being reported has expanded storage, elements unique to expanded storage are included in this section.

Note: Because of the larger memory models presented by z9xx mainframes and z/OS operating platforms, expanded storage is no longer needed. However, it is possible to configure central storage as expanded storage at IPL time to provide for compatability with previous platforms. Therefore, Performance Management will continue to support expanded storage measurements.

ONLINE ES The total number of online expanded storage
FRAMES: frames: PAGOLE.

AVAILABLE ES The average number of unused expanded storage
FRAMES: frames: PAGAVES.

AV PAGES TO The per-second rate of pages moved to expanded
ES RATE: storage: PAGAVPES. Use in conjunction with the
 migrated rate to analyze thrashing of expanded
 storage.

AV PAGES The per-second rate of pages moved (migrated)
MIGRATED from expanded storage to auxiliary storage:
RATE: PAGAVPEA.

AV ES FRAMES The rate of expanded storage frames being freed
FREED RATE: without being migrated: PAGPSFNM.

AV PRIM WS The average rate of primary working set
FR ES RATE: migration from expanded storage to auxiliary
 storage: PAGPSPWS. A primary working set is

the set of critical pages that must reside in central storage prior to an address space being dispatched. It consists of tables and queues, fixed pages, as well as the first page of each segment in the address space's working set.

AUX STORAGE SWAP RATE: The per-second rate of swaps to auxiliary storage data sets as derived by the following computation:

$$(PAGAXD+PAGLAX+PAGMIG) / DURATION$$

EXT STORAGE SWAP RATE: The per-second rate of swaps to expanded storage as derived by the following computation:

$$(PAGESED+PAGLES) / DURATION$$

PG RT FROM EXP STOR: The average number of page frames per second that were moved from expanded storage to real storage due to demand activity: PAGPSRES.

VIO PAGE RT TO EXP STOR: The rate per second at which VIO pages were transferred by the real storage manager (RSM) from central storage to expanded storage: PAGPSVME.

VIO MIGR RT FROM E STOR: The rate per second at which VIO pages were migrated from expanded storage through central storage to auxiliary storage: PAGPSVMG.

VIO PAGE RT FROM E STOR: The rate per second at which VIO pages were transferred by the real storage manager (RSM) from expanded storage to central storage: PAGPSVRE.

HIP PAGE RT TO EXP STOR: The rate per second at which hiperspace pages were transferred by the real storage manager (RSM) from central storage to expanded storage: PAGPSHME.

HIP MIGR RT FROM E STOR: The rate per second at which hiperspace pages were migrated from expanded storage through central storage to auxiliary storage: PAGPSHMG.

HIP PAGE RT FROM E STOR: The rate per second at which hiperspace pages were transferred by the real storage manager (RSM) from expanded storage to central storage: PAGPSHRE.

AV CSA IN The average number of CSA pages in expanded
EXP STOR: storage: PAGAVCSE.

AV LPA IN The average number of LPA pages in expanded
EXP STOR: storage: PAGAVLPE.

AV LSQA IN The average number of LSQA pages in expanded
EXP STOR: storage: PAGAVLQE.

AV REG+SWA The average number of Region plus SWA pages in
IN EXP STOR: expanded storage: PAGAVRGS.

AV SQA IN The average number of SQA pages in expanded
EXP STOR: storage: PAGAVSQE.

AV HIP IN The average number of hiperspace pages in
EXP STOR: expanded storage: PAGAVHPX.

AV VIO IN The average number of VIO pages in expanded
EXP STOR: storage: PAGAVVPX.

CA MICS Performance Manager
MONTHLY REAL STORAGE SYSTEM USAGE REPORT
SYSTEM = ESA1 RUNDATE = APR 4, yyyy

MONTHS (YY-MM):	91-02	91-03	CHANGE %

REAL STORAGE USE			

NUCLEUS FRAMES	686.00	686.00	---
PRIVATE FRAMES	2217.50	2279.37	2.79 %
LSQA FIXED FRAMES	481.06	485.63	0.95 %
CSA FRAMES	414.05	407.34	-1.62 %
LPA FRAMES	1056.17	1055.01	-0.10 %
SQA FIXED FRAMES	414.22	406.47	-1.87 %

REAL STORAGE AVAILABILITY			

FRAMES IN REAL STORAGE	5406.00	5406.00	-0.00 %
AVAILABLE FRAME COUNT	956.96	899.84	-5.96 %

FIXED STORAGE USE			

NUCLEUS FRAMES	686.00	686.00	---
PRIVATE FIXED FRAMES	100.29	109.77	9.45 %
CSA FIXED FRAMES	15.60	15.41	-1.22 %
LSQA FIXED FRAMES	481.06	485.63	0.95 %
SQA FIXED FRAMES	414.22	406.47	-1.87 %
LPA FIXED FRAMES	61.68	60.35	-2.15 %
FIXED BELOW 16 MB (XA)	545.50	557.61	2.22 %

FIXED STORAGE AVAILABILITY			

TOTAL FIXED FRAMES	1080.97	1104.54	2.18 %
TOTAL PAGEABLE FRAMES	5406.63	5406.12	-0.00 %

Figure 2-5. Sample Monthly Real Storage System Usage Report (Part 1 of 2)

CA MICS Performance Manager			
MONTHLY REAL STORAGE SYSTEM USAGE REPORT			
SYSTEM = ESA1 RUNDATE = APR 4, yyyy			

MONTHS (YY-MM):	91-02	91-03	CHANGE %

AUXILIARY STORAGE WORKLOAD			

DEMAND PAGE RATE	6.28	5.88	-6.32 %
LPA PAGES PER SECOND	0.42	0.42	---
CSA PAGES PER SECOND	0.90	0.90	---
MAX DEMAND PAGE RATE	10.46	9.97	-4.69 %
BLOCK IN RATE	2.06	1.09	-47.09 %
PAGES PER BLOCK IN	13.50	18.40	39.26 %
SWAP RATE	1.02	0.96	-5.88 %
MAX SWAP RATE	5.48	5.34	-2.55 %
PAGES PER SWAP IN	61.04	60.82	-0.36 %
TOTAL PAGES PER SECOND	61.00	51.22	-16.03 %
PAGES MOVED/SEC	9.11	9.07	-0.43 %

VIRTUAL:REAL RATIO	2.58	2.44	-5.42 %

EXPANDED STORAGE WORKLOAD			

ONLINE ES FRAMES	6144.00	6144.00	---
AVAILABLE ES FRAMES	3035.50	3120.71	2.81 %
AV PAGES TO ES RATE	1.90	1.72	-9.47 %
AV PAGES MIGRATED RATE	0.04	0.01	-75.00 %
AV ES FRAMES FREED RATE	0.02	0.01	-50.00 %
AV PRIM WS FR ES RATE	0.06	0.04	-33.33 %
AUX STORAGE SWAP RATE	0.00	0.00	N/A
EXT STORAGE SWAP RATE	1.02	0.96	-5.88 %
PG RT FROM EXP STOR	1.30	0.73	-43.85 %
VIO PAGE RT TO EXP STOR	0.00		N/A
VIO MIGR RT FROM E STOR	0.00		N/A
VIO PAGE RT FROM E STOR	0.00		N/A
HIP PAGE RT TO EXP STOR	0.55	0.47	-14.55 %
HIP MIGR RT FROM E STOR	0.02	0.02	---
HIP PAGE RT FROM E STOR	0.53	0.51	-3.77 %
AV CSA IN EXP STOR	31.50	35.87	13.87 %
AV LPA IN EXP STOR	1.50	1.61	7.33 %
AV LSQA IN EXP STOR	50.00	52.88	5.76 %
AV REG+SWA IN EXP STOR	3083.59	2971.66	-3.62 %
AV SQA IN EXP STOR	4.61	5.30	14.97 %
AV HIP IN EXP STOR	6.57	9.43	45.53 %
AV VIO IN EXP STOR	0.00		N/A

Figure 2-5. Sample Monthly Real Storage System Usage Report (Part 2 of 2)

CA MICS Performance Manager
 HOURLY REAL STORAGE SYSTEM USAGE REPORT
 SYSTEM = TA01 RUNDATE = SEP 8, yyyy

HOUR-DATE (HR-MM/DD):	09-09/06	10-09/06	11-09/06	12-09/06	13-09/06	14-09/06	CHANGE %
REAL STORAGE USE							
NUCLEUS FRAMES	533.00	533.00	533.00	533.00	533.00	533.00	---
PRIVATE FRAMES	4648.00	4663.50	4691.50	4601.75	4008.00	4294.25	-7.61 %
LSQA FIXED FRAMES	743.00	819.00	828.00	800.00	746.25	814.75	9.65 %
CSA FRAMES	561.00	563.00	522.50	435.00	446.75	447.50	-20.23 %
LPA FRAMES	1211.00	1237.50	1198.50	1161.50	1197.00	1219.25	0.68 %
SQA FIXED FRAMES	639.25	640.25	641.25	641.75	642.25	641.75	0.39 %
REAL STORAGE AVAILABILITY							
FRAMES IN REAL STORAGE	7536.00	7536.00	7536.00	7536.00	7536.00	7536.00	---
AVAILABLE FRAME COUNT	653.50	486.00	544.25	754.00	1253.75	887.00	35.73 %
FIXED STORAGE USE							
NUCLEUS FRAMES	533.00	533.00	533.00	533.00	533.00	533.00	---
PRIVATE FIXED FRAMES	810.75	764.50	781.50	750.50	650.50	660.25	-18.56 %
CSA FIXED FRAMES	54.00	54.00	54.00	53.50	53.50	53.75	-0.46 %
LSQA FIXED FRAMES	743.00	819.00	828.00	800.00	746.25	814.75	9.65 %
SQA FIXED FRAMES	639.25	640.25	641.25	641.75	642.25	641.75	0.39 %
LPA FIXED FRAMES	57.00	57.00	57.00	57.00	57.00	57.00	---
FIXED BELOW 16 MB (XA)	1254.25	1291.50	1288.25	1209.50	1097.25	1177.75	-6.09 %
FIXED STORAGE AVAILABILITY							
TOTAL FIXED FRAMES	2305.00	2335.25	2362.25	2303.50	2150.75	2228.75	-3.30 %
TOTAL PAGEABLE FRAMES	7536.00	7536.00	7536.00	7536.00	7536.00	7536.00	---

Figure 2-6. Sample Hourly Real Storage System Usage Report (Part 1 of 2)

CA MICS Performance Manager							
HOURLY REAL STORAGE SYSTEM USAGE REPORT							
SYSTEM = TA01 RUNDATE = SEP 8, yyyy							

HOURLY DATE (HR-MM/DD):	09-09/06	10-09/06	11-09/06	12-09/06	13-09/06	14-09/06	CHANGE %

AUXILIARY STORAGE WORKLOAD							

DEMAND PAGE RATE	2.00	2.85	2.48	2.22	1.49	1.74	-13.00 %
LPA PAGES PER SECOND	0.00	0.02	0.00	0.01	0.01	0.00	N/A
CSA PAGES PER SECOND	0.00	0.02	0.03	0.08	0.01	0.00	N/A
MAX DEMAND PAGE RATE	3.04	3.34	3.01	2.74	2.16	1.96	-35.52 %
BLOCK IN RATE	0.44	1.25	2.04	0.67	1.26	3.91	778.64 %
PAGES PER BLOCK IN	12.01	9.03	9.98	11.05	10.09	9.17	-23.65 %
SWAP RATE	0.23	0.29	0.30	0.30	0.28	0.37	60.86 %
MAX SWAP RATE	0.30	0.33	0.32	0.33	0.34	0.42	40.00 %
PAGES PER SWAP IN	79.83	77.56	73.65	76.20	67.95	70.14	-12.13 %
TOTAL PAGES PER SECOND	44.43	60.59	68.39	56.43	52.75	90.26	103.15 %
PAGES MOVED/SEC	2.62	2.67	3.92	3.78	3.87	4.17	59.16 %
VIRTUAL:REAL RATIO	1.81	1.92	1.93	1.96	2.23	2.13	17.67 %

Figure 2-6. Sample Hourly Real Storage System Usage Report (Part 2 of 2)

2.2.3.3 Real Storage Private Usage Report

The Real Storage Private Usage Report tracks the use of real storage frames at the service class or service class period level. Both service and report classes are available for reporting. The Real Storage Private Usage Report can be produced at any CA MICS timespan, and can report on any numeric element present in the selected input file. All selected data is processed by a single report request. Data is always sorted for separate presentation by SYSID and can also be reported by ZONE and HOUR if you select certain timespans.

REPORT OPTIONS

A separate report is produced for each SYSID you select. The Real Storage Private Usage Report supports standard time selection options as well SYSID selection and external data selection. These options are discussed in Section 2.2.7.

DATA ELEMENTS

The Real Storage Private Usage Report shows the average number of frames occupied by each transaction in each reporting group during the reported time period. This number does not include SQA, LPA, and CSA frames used (LSQA is included). The value is taken directly from the data element measuring average total frames in both central and expanded storage (fffAVAFt) for the selected CA MICS file. This element name is the standard (default) value for the report. You can optionally select any single meaningful numeric data element from any supported input file (WLMSEC or WLM_EC) for display and analysis in this report. Consult the CA MICS data dictionary for a complete list of the elements that you can select. (Prior to RMF 3.5, fffAVRSF can be used for analyzing storage usage instead of fffAVAFt, which first became available with RMF 3.5.)

Figure 2-7 shows an example of the Real Storage Private Usage Report at the weekly level. The last column represents the percentage of change in real storage use from the base time period (taken to be the first period containing data for the service class). If there is data for only one period, dashes (---) are printed for the percent of change. If the base value is zero, N/A is printed to indicate the percent changed cannot be calculated.

CA MICS Performance Manager HOURLY REAL STORAGE "PRIVATE" USAGE REPORT FOR ZONE = 1 SYSTEM = TS04 RUNDATA = JAN 12, yyyy REPORTING ON SERVICE CLASS ACTIVITY DATA										
HOUR/WEEK (HR-MM/DD):	08-01/11	09-01/11	10-01/11	11-01/11	12-01/11	13-01/11	14-01/11	15-01/11	16-01/11	CHANGE %
AVERAGE REAL FRAMES										
SYSTEM TASKS	315.03	317.08	317.41	314.98	306.80	307.42	308.03	308.59	313.76	-0.40 %
TSO EXTERNAL USERS	869.93	883.20	866.90	924.86	930.69	939.50	942.00	941.62	942.00	8.28 %
BATCH TESTING	174.72	183.56	166.89	176.15	200.46	170.02	166.87	164.20	182.31	4.34 %
STC - HIGH PRIORITY	290.91	199.15	249.74	206.98	203.19	229.18	192.83	203.80	275.67	-5.23 %
TSO LOW PRIORITY	188.90	204.64	164.76	171.97	176.87	210.89	298.92	273.80	154.79	-18.05 %
DISCRETIONARY	266.79	281.96	227.85	227.05	435.54	268.72	290.06	198.34	476.36	78.55 %
SPECIAL PROCESSING	209.89	171.11	160.59	190.89	120.75	128.51	247.86	239.93	233.64	11.31 %
TSO INTERNAL USERS	1176.51	1214.02	1202.94	1168.08	1168.13	1197.75	1225.29	1214.07	1071.77	-8.90 %
STC - MEDIUM PRIORITY	763.65	785.14	762.20	764.86	790.98	792.17	785.91	772.73	761.18	-0.32 %
STC - LOW PRIORITY	181.85	177.40	191.94	197.44	199.76	205.28	208.27	203.94	200.16	10.06 %
GROUP 53	81.52	81.75	81.73	81.69	81.69	81.70	81.70	81.70	81.69	0.20 %
BATCH PRODUCTION	1265.93	1070.42	1021.61	1031.10	1038.21	1053.70	1058.86	1067.63	1070.39	-15.44 %
JES2	1116.39	945.41	967.77	1073.90	1134.45	1147.78	1163.18	1156.53	1156.06	3.55 %
NETWORK SERVICES	1048.75	1049.05	1002.59	743.68	635.02	635.04	635.06	608.26	484.06	-53.84 %

Figure 2-7. Sample Hourly Private Storage Usage Report

2.2.3.4 Real Storage Swap Reason Report

RMF versions at level 3.3 or above provide enhanced swapping statistics. This data is maintained in the CA MICS SCPSWP file. The data provides swap counts indicating the cause or reason for the swap and the type or placement of the swap (real memory, extended storage, or auxiliary storage). Analysis of this information can provide you with insight into the causes of swapping overhead, which you can then address to improve system performance.

The Real Storage Swap Reason Report lists the number of different types of swaps that occurred by the reason for the swap. Figure 2-8 shows a sample Real Storage Swap Reason Report.

The following swap reasons are reported:

TERMINAL OUTPUT WAIT:	Swaps caused by terminal output buffers being exhausted. See the element PAGTOS in the CA MICS data dictionary.
TERMINAL INPUT WAIT:	Swaps caused by a terminal waiting for input. See the element PAGTIS in the CA MICS Data dictionary.
LONG WAIT:	Swaps caused by a task entering a long wait condition. See the element PAGLWS in the CA MICS data dictionary.
AUX. STORAGE SHORTAGE:	Swaps caused by a shortage of auxiliary storage. See the element PAGASS in the CA MICS data dictionary.
REAL STORAGE SHORTAGE:	Swaps caused by a shortage of real storage. See the element PAGRSS in the CA MICS data dictionary.
DETECTED WAIT:	Swaps caused by the SRM detecting that a task is waiting. See the element PAGDWS in the CA MICS data dictionary.
REQUESTED:	Swaps caused by a task requesting its pages be moved to real storage. See the element PAGNRS in the CA MICS data dictionary.
ENQUEUE EXCHANGE:	Swaps done by the SRM to correct an enqueue problem. See the element PAGEES in the CA MICS data dictionary.

EXCHANGE ON REC.	Swaps caused when the number of ready tasks in a period exceeds the multiprogramming level. See the element PAGEXS in the CA MICS data dictionary.
UNILATERAL:	Swaps made by the SRM because the target multiprogramming level for a period has been exceeded. See the element PAGULS in the CA MICS data dictionary.
TRANS. TO NON-SWAP:	Swaps made because a task requested that it be made non-swappable. See the element PAGTXS in the CA MICS data dictionary.
IMPROVE SYST. PAGING RATES:	Swaps performed in order to improve the system paging rate. See the element PAGISP in the CA MICS data dictionary.
IMPROVE CEN. STG. USE:	Swaps performed in order to improve central storage usage. See the element PAGISU in the CA MICS data dictionary.
ROOM FOR LONG- SWAPPED:	Swaps performed in order to make room in central storage to swap in an address space that has been swapped out too long. See the element PAGRLS in the CA MICS data dictionary.
APPC WAIT:	Swaps performed due to APPC conversation wait conditions. The SRM may choose to swap an address space that is waiting for APPC services. See the element PAGAPC in the CA MICS data dictionary.
OMVS INPUT WAIT:	Swaps caused by OpenEdition MVS waiting for input.
OMVS OUTPUT WAIT:	Swaps caused by output wait conditions for OpenEdition MVS.

CA MICS Performance Manager										
HOURLY REAL STORAGE "SWAP-REASON" REPORT										
SYSTEM = PROD RUNDTE = Mar 7, yyyy										
16:51 Monday, March 7, yyyy 1										
HOURLY-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %

TERMINAL OUTPUT WAIT										

TOTAL SWAP CANDIDATES	72.00	44.00	38.00	24.00	44.00	9.00	16.00	22.00	58.00	-19.44 %
TOTAL LOGICAL SWAPS	69.00	44.00	36.00	24.00	42.00	9.00	14.00	17.00	50.00	-27.53 %
EFFECTIVE LOGICAL SWAPS	65.00	44.00	34.00	23.00	40.00	9.00	14.00	17.00	45.00	-30.76 %
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	7.00	0.00	4.00	1.00	4.00	0.00	1.00	5.00	13.00	85.71 %
EXT STOR DIR LOG. SWAPS	4.00	0.00	2.00	1.00	2.00	0.00	0.00	0.00	5.00	25.00 %
EXT STOR DIR PHYS SWAPS	3.00	0.00	2.00	0.00	2.00	0.00	1.00	5.00	8.00	166.66 %
EFFECTIVE LOG/EXT SWAPS	72.00	44.00	38.00	24.00	44.00	9.00	15.00	22.00	58.00	-19.44 %

TERMINAL INPUT WAIT										

TOTAL SWAP CANDIDATES	542.00	332.00	499.00	665.00	880.00	611.00	445.00	231.00	425.00	-21.58 %
TOTAL LOGICAL SWAPS	520.00	321.00	449.00	602.00	815.00	571.00	414.00	184.00	354.00	-31.92 %
EFFECTIVE LOGICAL SWAPS	432.00	272.00	375.00	452.00	652.00	441.00	303.00	132.00	242.00	-43.98 %
TOTAL AUX STORAGE SWAPS	0.00	1.00	13.00	40.00	14.00	2.00	46.00	0.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	6.00	2.00	0.00	7.00	0.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	24.00	7.00	0.00	24.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	1.00	13.00	10.00	5.00	2.00	15.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	110.00	60.00	124.00	183.00	219.00	170.00	111.00	99.00	183.00	66.36 %
EXT STOR DIR LOG. SWAPS	88.00	49.00	74.00	126.00	156.00	130.00	87.00	52.00	112.00	27.27 %
EXT STOR DIR PHYS SWAPS	22.00	11.00	50.00	57.00	63.00	40.00	24.00	47.00	71.00	222.72 %
EFFECTIVE LOG/EXT SWAPS	542.00	331.00	486.00	625.00	866.00	609.00	399.00	231.00	425.00	-21.58 %

Figure 2-8. Sample Real Storage Swap Reason Report (Part 1 of 9)

16:51 Monday, March 7, yyyy 2

CA MICS Performance Manager
HOURLY REAL STORAGE "SWAP-REASON" REPORT
SYSTEM = PROD RUNDATE = Mar 7, yyyy

HOUR-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %

LONG WAIT										

TOTAL SWAP CANDIDATES	54.00	178.00	48.00	27.00	122.00	81.00	120.00	42.00	38.00	-29.62 %
TOTAL LOGICAL SWAPS	49.00	169.00	38.00	18.00	112.00	77.00	115.00	26.00	26.00	-46.93 %
EFFECTIVE LOGICAL SWAPS	36.00	160.00	27.00	10.00	87.00	63.00	83.00	14.00	15.00	-58.33 %
TOTAL AUX STORAGE SWAPS	0.00	1.00	2.00	6.00	4.00	0.00	25.00	0.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	1.00	2.00	0.00	3.00	0.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	1.00	0.00	3.00	1.00	0.00	17.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	0.00	2.00	2.00	1.00	0.00	5.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	18.00	17.00	21.00	13.00	32.00	18.00	17.00	28.00	23.00	27.77 %
EXT STOR DIR LOG. SWAPS	13.00	8.00	11.00	5.00	24.00	14.00	15.00	12.00	11.00	-15.38 %
EXT STOR DIR PHYS SWAPS	5.00	9.00	10.00	8.00	8.00	4.00	2.00	16.00	12.00	140.00 %
EFFECTIVE LOG/EXT SWAPS	54.00	177.00	46.00	21.00	118.00	81.00	95.00	42.00	38.00	-29.62 %

AUX. STORAGE SHORTAGE										

TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-8. Sample Real Storage Swap Reason Report (Part 2 of 9)

CA MICS Performance Manager										16:51 Monday, March 7, yyyy	3
HOURLY REAL STORAGE "SWAP-REASON" REPORT											
SYSTEM = PROD RUNDATE = Mar 7, yyyy											
HOURLY-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE	%
REAL STORAGE SHORTAGE											
TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
DETECTED WAIT											
TOTAL SWAP CANDIDATES	133.00	452.00	559.00	654.00	681.00	177.00	131.00	119.00	124.00	-6.76	%
TOTAL LOGICAL SWAPS	119.00	430.00	517.00	589.00	590.00	158.00	96.00	64.00	80.00	-32.77	%
EFFECTIVE LOGICAL SWAPS	78.00	375.00	433.00	414.00	357.00	93.00	42.00	28.00	17.00	-78.20	%
TOTAL AUX STORAGE SWAPS	0.00	0.00	2.00	8.00	2.00	1.00	19.00	1.00	0.00	N/A	
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	1.00	0.00	1.00	8.00	1.00	0.00	N/A	
AUX STOR DIR LOG. SWAPS	0.00	0.00	1.00	7.00	2.00	0.00	10.00	0.00	0.00	N/A	
SWAPS MIGRATED TO AUX.	0.00	0.00	1.00	0.00	0.00	0.00	1.00	0.00	0.00	N/A	
TOTAL EXT STORAGE SWAPS	55.00	77.00	125.00	232.00	322.00	83.00	71.00	90.00	107.00	94.54	%
EXT STOR DIR LOG. SWAPS	41.00	55.00	83.00	168.00	231.00	65.00	44.00	36.00	63.00	53.65	%
EXT STOR DIR PHYS SWAPS	14.00	22.00	42.00	64.00	91.00	18.00	27.00	54.00	44.00	214.28	%
EFFECTIVE LOG/EXT SWAPS	133.00	452.00	557.00	646.00	679.00	176.00	112.00	118.00	124.00	-6.76	%

Figure 2-8. Sample Real Storage Swap Reason Report (Part 3 of 9)

16:51 Monday, March 7, yyyy 4

CA MICS Performance Manager
 HOURLY REAL STORAGE "SWAP-REASON" REPORT
 SYSTEM = PROD RUNDATE = Mar 7, yyyy

HOURLY-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE	%
REQUESTED											
TOTAL SWAP CANDIDATES	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00	%
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL EXT STORAGE SWAPS	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00	%
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR PHYS SWAPS	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00	%
EFFECTIVE LOG/EXT SWAPS	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00	%
ENQUEUE EXCHANGE											
TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	

REQUESTED

TOTAL SWAP CANDIDATES	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00	%
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL EXT STORAGE SWAPS	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00	%
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR PHYS SWAPS	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00	%
EFFECTIVE LOG/EXT SWAPS	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00	%

ENQUEUE EXCHANGE

TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	

Figure 2-8. Sample Real Storage Swap Reason Report (Part 4 of 9)

CA MICS Performance Manager										16:51 Monday, March 7, yyyy 5
HOURLY REAL STORAGE "SWAP-REASON" REPORT										
SYSTEM = PROD RUNDATE = Mar 7, yyyy										
HOURLY-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %
EXCHANGE ON REC.										
TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
UNILATERAL										
TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-8. Sample Real Storage Swap Reason Report (Part 5 of 9)

CA MICS Performance Manager										
HOURLY REAL STORAGE "SWAP-REASON" REPORT										
SYSTEM = PROD RUNDATE = Mar 7, yyyy										

HOURLY-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %

TRANS. TO NON-SWAP.										

TOTAL SWAP CANDIDATES	4.00	4.00	14.00	8.00	15.00	7.00	14.00	9.00	3.00	-25.00 %
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	2.00	0.00	1.00	4.00	1.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	2.00	0.00	1.00	4.00	1.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	4.00	4.00	14.00	6.00	15.00	6.00	10.00	8.00	3.00	-25.00 %
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR PHYS SWAPS	4.00	4.00	14.00	6.00	15.00	6.00	10.00	8.00	3.00	-25.00 %
EFFECTIVE LOG/EXT SWAPS	4.00	4.00	14.00	6.00	15.00	6.00	10.00	8.00	3.00	-25.00 %

IMPROVE CEN. STG. USE										

TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-8. Sample Real Storage Swap Reason Report (Part 6 of 9)

CA MICS Performance Manager										16:51 Monday, March 7, yyyy 7
HOURLY REAL STORAGE "SWAP-REASON" REPORT										
SYSTEM = PROD RUNDATE = Mar 7, yyyy										
HOURLY-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %
IMPROVE SYST. PAGING										
TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ROOM FOR LONG-SWAPPED										
TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-8. Sample Real Storage Swap Reason Report (Part 7 of 9)

CA MICS Performance Manager										16:51 Monday, March 7, yyyy	8
HOURLY REAL STORAGE "SWAP-REASON" REPORT											
SYSTEM = PROD RUNDATE = Mar 7, yyyy											
HOURLY-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE	%

APPC WAIT											

TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	

OMVS INPUT WAIT											

TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	

Figure 2-8. Sample Real Storage Swap Reason Report (Part 8 of 9)

CA MICS Performance Manager										16:51 Monday, March 7, yyyy	9
HOURLY REAL STORAGE "SWAP-REASON" REPORT											
SYSTEM = PROD RUNDATE = Mar 7, yyyy											

HOURLY-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE	%

OMVS OUTPUT WAIT											

TOTAL SWAP CANDIDATES	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOGICAL SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL AUX STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
AUX STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
SWAPS MIGRATED TO AUX.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TOTAL EXT STORAGE SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR LOG. SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXT STOR DIR PHYS SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EFFECTIVE LOG/EXT SWAPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	

Figure 2-8. Sample Real Storage Swap Reason Report (Part 9 of 9)

2.2.3.5 Real Storage Swap Type Report

RMF versions at level 3.3 or above provide enhanced swapping statistics. This data is maintained in the CA MICS SCPSWP File. It provides swap counts indicating the cause or reason for the swap and the type or placement of the swap (central storage, extended storage, or auxiliary storage). Analysis of this information can provide you with insight into the causes of swapping overhead, which you can then address to improve system performance.

The Real Storage Swap Type Report lists the number of different reasons for the swaps that occurred by the type of the swap. Figure 2-9 shows a sample Real Storage Swap Type Report.

The following swap types are reported:

TOTAL SWAPS:	The total number of swap operations. See the element SWPTOSWP in the CA MICS Data Encyclopedia.
TOTAL LOGICAL SWAPS:	The total number of logical swap operations. See the element SWPLSWPS in the CA MICS Data Encyclopedia.
EFFECTIVE LOGICAL SWAPS:	The total number of effective logical swaps. See the element SWPLSWPE in the CA MICS Data Encyclopedia.
TOTAL AUX STORAGE SWAPS:	The total number of swaps directed towards auxiliary storage. See the element SWPAUX in the CA MICS Data Encyclopedia.
AUX STOR DIR PHYS SWAPS:	The total number of physical swaps directed towards auxiliary storage. See the element SWPAXD in the CA MICS Data Encyclopedia.
AUX STOR DIR LOG SWAPS:	The total number of logical swaps directed towards auxiliary storage. See the element SWPLAX in the CA MICS Data Encyclopedia.
SWAPS MIGRATED TO AUX:	The total number of swaps migrated from extended storage to auxiliary storage. See the element SWPMIG in the CA MICS Data Encyclopedia.
TOTAL EXT STORAGE SWAPS:	The total number of swaps directed towards extended storage. See the element SWPES in the CA MICS Data Encyclopedia.

EXT STOR The total number of logical swaps directed
 DIR LOG towards extended storage. See the element
 SWAPS: SWPLES in the CA MICS Data Encyclopedia.

EXT STOR The total number of physical swaps directed
 DIR PHYS towards extended storage. See the element
 SWAPS: SWPESD in the CA MICS Data Encyclopedia.

EFFECTIVE The total number of effective logical extended
 LOG/EXT storage swaps. See the element SWPNSWP in the
 SWAPS: CA MICS Data Encyclopedia.

NOTE: Some of the elements listed above appear on your
 reports only if you have extended storage on your processor
 and are using the correct version of RMF.

CA MICS Performance Manager										
HOURLY REAL STORAGE "SWAP-TYPE" REPORT										
SYSTEM = PROD RUNDATE = Mar 7, yyyy										

HOUR-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %

TOTAL SWAP CANDIDATES										

TERMINAL OUTPUT WAIT	72.00	44.00	38.00	24.00	44.00	9.00	16.00	22.00	58.00	-19.44 %
TERMINAL INPUT WAIT	542.00	332.00	499.00	665.00	880.00	611.00	445.00	231.00	425.00	-21.58 %
LONG WAIT	54.00	178.00	48.00	27.00	122.00	81.00	120.00	42.00	38.00	-29.62 %
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
DETECTED WAIT	133.00	452.00	559.00	654.00	681.00	177.00	131.00	119.00	124.00	-6.76 %
REQUESTED	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00 %
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TRANS. TO NON-SWAP.	4.00	4.00	14.00	8.00	15.00	7.00	14.00	9.00	3.00	-25.00 %
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-9. Sample Real Storage Swap Type Report (Part 1 of 11)

CA MICS Performance Manager										
HOURLY REAL STORAGE "SWAP-TYPE" REPORT										
SYSTEM = PROD RUNDATE = Mar 7, yyyy										

HOUR-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %

TOTAL LOGICAL SWAPS										

TERMINAL OUTPUT WAIT	69.00	44.00	36.00	24.00	42.00	9.00	14.00	17.00	50.00	-27.53 %
TERMINAL INPUT WAIT	520.00	321.00	449.00	602.00	815.00	571.00	414.00	184.00	354.00	-31.92 %
LONG WAIT	49.00	169.00	38.00	18.00	112.00	77.00	115.00	26.00	26.00	-46.93 %
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
DETECTED WAIT	119.00	430.00	517.00	589.00	590.00	158.00	96.00	64.00	80.00	-32.77 %
REQUESTED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TRANS. TO NON-SWAP.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-9. Sample Real Storage Swap Type Report (Part 2 of 11)

CA MICS Performance Manager HOURLY REAL STORAGE "SWAP-TYPE" REPORT SYSTEM = PROD RUNDATE = Mar 7, yyyy										
17:00 Monday, March 7, yyyy 3										

HOUR-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %

EFFECTIVE LOGICAL SWAPS										

TERMINAL OUTPUT WAIT	65.00	44.00	34.00	23.00	40.00	9.00	14.00	17.00	45.00	-30.76 %
TERMINAL INPUT WAIT	432.00	272.00	375.00	452.00	652.00	441.00	303.00	132.00	242.00	-43.98 %
LONG WAIT	36.00	160.00	27.00	10.00	87.00	63.00	83.00	14.00	15.00	-58.33 %
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
DETECTED WAIT	78.00	375.00	433.00	414.00	357.00	93.00	42.00	28.00	17.00	-78.20 %
REQUESTED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TRANS. TO NON-SWAP.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-9. Sample Real Storage Swap Type Report (Part 3 of 11)

CA MICS Performance Manager										17:00 Monday, March 7, yyyy	4
HOURLY REAL STORAGE "SWAP-TYPE" REPORT											
SYSTEM = PROD RUNDATE = Mar 7, yyyy											
HOURLY DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE	%

TOTAL AUX STORAGE SWAPS											

TERMINAL OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	N/A	
TERMINAL INPUT WAIT	0.00	1.00	13.00	40.00	14.00	2.00	46.00	0.00	0.00	N/A	
LONG WAIT	0.00	1.00	2.00	6.00	4.00	0.00	25.00	0.00	0.00	N/A	
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
DETECTED WAIT	0.00	0.00	2.00	8.00	2.00	1.00	19.00	1.00	0.00	N/A	
REQUESTED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TRANS. TO NON-SWAP.	0.00	0.00	0.00	2.00	0.00	1.00	4.00	1.00	0.00	N/A	
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	

Figure 2-9. Sample Real Storage Swap Type Report (Part 4 of 11)

CA MICS Performance Manager										17:00 Monday, March 7, yyyy	5
HOURLY REAL STORAGE "SWAP-TYPE" REPORT											
SYSTEM = PROD RUNDATE = Mar 7, yyyy											
HOURLY-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE	%
AUX STOR DIR PHYS SWAPS											
TERMINAL OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	N/A	
TERMINAL INPUT WAIT	0.00	0.00	0.00	6.00	2.00	0.00	7.00	0.00	0.00	N/A	
LONG WAIT	0.00	0.00	0.00	1.00	2.00	0.00	3.00	0.00	0.00	N/A	
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
DETECTED WAIT	0.00	0.00	0.00	1.00	0.00	1.00	8.00	1.00	0.00	N/A	
REQUESTED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TRANS. TO NON-SWAP.	0.00	0.00	0.00	2.00	0.00	1.00	4.00	1.00	0.00	N/A	
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	

Figure 2-9. Sample Real Storage Swap Type Report (Part 5 of 11)

17:00 Monday, March 7, yyyy6

CA MICS Performance Manager

HOURLY REAL STORAGE "SWAP-TYPE" REPORT

SYSTEM = PROD RUNDATE = Mar 7, yyyy

HOUR-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE	%

AUX STOR DIR LOG. SWAPS											

TERMINAL OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TERMINAL INPUT WAIT	0.00	0.00	0.00	24.00	7.00	0.00	24.00	0.00	0.00	N/A	
LONG WAIT	0.00	1.00	0.00	3.00	1.00	0.00	17.00	0.00	0.00	N/A	
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
DETECTED WAIT	0.00	0.00	1.00	7.00	2.00	0.00	10.00	0.00	0.00	N/A	
REQUESTED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TRANS. TO NON-SWAP.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	

Figure 2-9. Sample Real Storage Swap Type Report (Part 6 of 11)

CA MICS Performance Manager										17:00 Monday, March 7, yyyy 7
HOURLY REAL STORAGE "SWAP-TYPE" REPORT										
SYSTEM = PROD RUNDATE = Mar 7, yyyy										
HOURLY DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %
SWAPS MIGRATED TO AUX.										
TERMINAL OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TERMINAL INPUT WAIT	0.00	1.00	13.00	10.00	5.00	2.00	15.00	0.00	0.00	N/A
LONG WAIT	0.00	0.00	2.00	2.00	1.00	0.00	5.00	0.00	0.00	N/A
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
DETECTED WAIT	0.00	0.00	1.00	0.00	0.00	0.00	1.00	0.00	0.00	N/A
REQUESTED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TRANS. TO NON-SWAP.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-9. Sample Real Storage Swap Type Report (Part 7 of 11)

CA MICS Performance Manager										17:00 Monday, March 7, yyyy	8
HOURLY REAL STORAGE "SWAP-TYPE" REPORT											
SYSTEM = PROD RUNDATE = Mar 7, yyyy											
HOURLY DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE	%

TOTAL EXT STORAGE SWAPS											

TERMINAL OUTPUT WAIT	7.00	0.00	4.00	1.00	4.00	0.00	1.00	5.00	13.00	85.71	%
TERMINAL INPUT WAIT	110.00	60.00	124.00	183.00	219.00	170.00	111.00	99.00	183.00	66.36	%
LONG WAIT	18.00	17.00	21.00	13.00	32.00	18.00	17.00	28.00	23.00	27.77	%
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
DETECTED WAIT	55.00	77.00	125.00	232.00	322.00	83.00	71.00	90.00	107.00	94.54	%
REQUESTED	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00	%
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
TRANS. TO NON-SWAP.	4.00	4.00	14.00	6.00	15.00	6.00	10.00	8.00	3.00	-25.00	%
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A	

Figure 2-9. Sample Real Storage Swap Type Report (Part 8 of 11)

CA MICS Performance Manager HOURLY REAL STORAGE "SWAP-TYPE" REPORT SYSTEM = PROD RUNDATE = Mar 7, yyyy										
17:00 Monday, March 7, yyyy 9										
HOURLY DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %
EXT STOR DIR LOG. SWAPS										
TERMINAL OUTPUT WAIT	4.00	0.00	2.00	1.00	2.00	0.00	0.00	0.00	5.00	25.00 %
TERMINAL INPUT WAIT	88.00	49.00	74.00	126.00	156.00	130.00	87.00	52.00	112.00	27.27 %
LONG WAIT	13.00	8.00	11.00	5.00	24.00	14.00	15.00	12.00	11.00	-15.38 %
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
DETECTED WAIT	41.00	55.00	83.00	168.00	231.00	65.00	44.00	36.00	63.00	53.65 %
REQUESTED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TRANS. TO NON-SWAP.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-9. Sample Real Storage Swap Type Report (Part 9 of 11)

CA MICS Performance Manager HOURLY REAL STORAGE "SWAP-TYPE" REPORT SYSTEM = PROD RUNDATE = Mar 7, yyyy										
17:00 Monday, March 7, yyyy 10										
HOUR-DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %
EXT STOR DIR PHYS SWAPS										
TERMINAL OUTPUT WAIT	3.00	0.00	2.00	0.00	2.00	0.00	1.00	5.00	8.00	166.66 %
TERMINAL INPUT WAIT	22.00	11.00	50.00	57.00	63.00	40.00	24.00	47.00	71.00	222.72 %
LONG WAIT	5.00	9.00	10.00	8.00	8.00	4.00	2.00	16.00	12.00	140.00 %
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
DETECTED WAIT	14.00	22.00	42.00	64.00	91.00	18.00	27.00	54.00	44.00	214.28 %
REQUESTED	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00 %
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TRANS. TO NON-SWAP.	4.00	4.00	14.00	6.00	15.00	6.00	10.00	8.00	3.00	-25.00 %
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-9. Sample Real Storage Swap Type Report (Part 10 of 11)

CA MICS Performance Manager HOURLY REAL STORAGE "SWAP-TYPE" REPORT SYSTEM = PROD RUNDATE = Mar 7, yyyy										
17:00 Monday, March 7, yyyy 11										
HOURLY DATE (HR-MM/DD):	08-02/28	09-02/28	10-02/28	11-02/28	12-02/28	13-02/28	14-02/28	15-02/28	16-02/28	CHANGE %
EFFECTIVE LOG/EXT SWAPS										
TERMINAL OUTPUT WAIT	72.00	44.00	38.00	24.00	44.00	9.00	15.00	22.00	58.00	-19.44 %
TERMINAL INPUT WAIT	542.00	331.00	486.00	625.00	866.00	609.00	399.00	231.00	425.00	-21.58 %
LONG WAIT	54.00	177.00	46.00	21.00	118.00	81.00	95.00	42.00	38.00	-29.62 %
AUX. STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
REAL STORAGE SHORTAGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
DETECTED WAIT	133.00	452.00	557.00	646.00	679.00	176.00	112.00	118.00	124.00	-6.76 %
REQUESTED	1.00	0.00	0.00	0.00	1.00	1.00	3.00	1.00	2.00	100.00 %
ENQUEUE EXCHANGE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
EXCHANGE ON REC.	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
UNILATERAL	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TRANS. TO NON-SWAP.	4.00	4.00	14.00	6.00	15.00	6.00	10.00	8.00	3.00	-25.00 %
IMPROVE CEN. STG. USE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
IMPROVE SYST. PAGING	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
ROOM FOR LONG-SWAPPED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
APPC WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS INPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
OMVS OUTPUT WAIT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-9. Sample Real Storage Swap Type Report (Part 11 of 11)

2.2.4 MVS Concepts

The methodology section is an introduction to the concepts underlying Real Storage Management (RSM) and the control mechanisms available within MVS that affect real storage. Real storage management is the efficient allocation of the system's real memory to each of the various active programs in the system.

This section explains the concept of virtual storage as it relates to MVS systems and Real Storage Management. First, as a foundation for these concepts, we present a very brief history of the development of hardware addressing schemes leading up to the current virtual systems. Next, we describe the layout of the virtual storage in MVS. We then explain the parts of the virtual memory that can be controlled or tracked in terms of real storage management. Finally, an overview is provided of expanded storage and how this can be managed. We provide examples to illustrate these concepts. The following sections include discussions of each of these topics:

- 1 - Virtual Storage Concepts
- 2 - MVS Virtual Storage Overview
- 3 - MVS Real Storage Control Mechanisms
- 4 - Expanded Storage Concepts

2.2.4.1 Virtual Storage Concepts

This section introduces concepts common to virtual memory systems. A virtual storage system (or virtual memory system) is a system where executing (active) programs reside in auxiliary storage (i.e., on disk) and selected portions of the program are placed in real storage as needed.

The need for memory management arises from the development of hardware and software architecture to support multiprocessing and time-sharing systems. The earliest computer systems ran one job at a time in serial order. Thus, there was no sharing of resources. As computers and peripheral devices became faster, the desire to use these resources more efficiently led to the introduction of multiprogramming systems. A multiprogramming system allows several jobs to be in memory concurrently. This allows one job to use CPU resources while another is waiting on I/O and vice versa.

The next level of development occurred as people costs started to approach hardware costs. The desire to increase programmer productivity led to the introduction of time-sharing systems where many users shared memory for short periods of time. An interactive time-sharing system gave individual programmers access to computing power through terminals when they needed it. When a given terminal was not active, the system could occupy itself with other users or other tasks.

The development of virtual systems stemmed from the creation of these systems. Memory was a very expensive resource and the memory demands of the time-sharing systems greatly exceeded the capacity of the existing computer systems. This began the attempts to deal with the issue of real storage management.

The type of address resolution used by the hardware was the major factor determining what type of memory sharing scheme was used for managing real storage. The three types of addressing used by the various systems and some of the design implications for sharing memory are briefly discussed in Section 2.2.4.1.1. The purpose of describing these early systems is to develop the logical progression of these concepts for the reader not familiar with virtual systems.

Two critical issues arise from the time-sharing of real memory that are only mentioned in this chapter. They are data security and system integrity. Early systems largely ignored these subjects. In all time-shared systems, the real

storage management methodology must have data security as a primary goal.

The following sections discuss paging systems to give you a foundation for the discussion of the MVS paging algorithms in Section 2.2.4.1.2 and for the memory management schemes in Section 2.2.4.3.

- 1 - Address Resolution Schemes
- 2 - Paging Concepts

2.2.4.1.1 Address Resolution Schemes

The resolution of addresses is a central issue to virtual systems for two reasons. First, memory may not be shared unless the code and data of one user is protected against unauthorized modification by the other users (be it accidental or malicious). Second, the address resolution mechanism largely dictates the manner in which users may be moved in and out of real storage.

If a virtual system is to provide more apparent (virtual) storage than is actually available, the system must be able to allocate the real memory it has in an efficient manner. One method of moving users in and out of real storage is paging. This is discussed in Section 2.2.4.1.2.

This section briefly traces the development of address resolution by the IBM operating systems in order to provide a conceptual basis for discussing MVS real storage issues.

In general, there are three kinds of address resolution used:

- o Absolute addressing
- o Relative addressing
- o Dynamic addressing

There are three times in the life of a program when addresses may be assigned: when the program is written or compiled (absolute addressing), when the program is loaded into memory (relative addressing), and when the program is executed (dynamic addressing). Of these schemes, only dynamic addressing requires hardware to support the resolution of an address. In this section, we discuss each of these addressing schemes in terms of the problems they impose or solve for a virtual memory system.

ABSOLUTE ADDRESSING

The early computers did not have any hardware to support memory sharing, such as memory protection features or dynamic address resolution. These computers had an instruction format something like that shown in Figure 2-19:




Figure 2-19. Computer Instruction Format

where the instruction had an op-code (o) and a general register specification (x), and the remainder of the bits comprised the address (a). The amount of real memory was limited by the number of bits in the address portion of the instruction format. Programs were prevented from accessing or altering the data of other programs because real memory was not shared concurrently. Only one task was present in memory at one time. There was no effective protection for the resident operating system code.

In terms of memory management, the only virtual memory allocation scheme possible for implementing a time-sharing system on these machines was swapping. Swapping involves moving an entire program into real storage from auxiliary storage in order to make it active, or moving an active program out of real storage into auxiliary storage, making it inactive.

RELATIVE ADDRESSING (ADDRESS RESOLUTION AT PROGRAM LOAD)

Nothing in the instruction format shown in Figure 2-19 precluded address resolution at the time the program was loaded. However, without memory protection facilities such as the protect keys used by the IBM S/360 series, there was no reason to develop the program loaders and linkage editors to implement load time address resolution. In the IBM S/360 series, there are two levels of address constants and the instruction format is designed to make load time address resolution more efficient. The S/360 instruction formation shown in Figure 2-20 provides two other advantages for program load address resolution.

```
oorxbaaa
|  |  |  |
|  |  |  |
|  |  |  |  +----- address offsets (12 bits)
|  |  |
|  |  |  +----- Base Register (4 bits)
|  |
|  |  +----- Index register (4 bits)
|  |
|  |  +----- General Register (4 bits)
|
+----- Instruction op-code (8 bits)
```

Figure 2-20. S/360 RX Instruction Format

The instruction length of an RX instruction is 32 bits (one full word). The address field of the instruction is defined by the baaa portion, which is 16 bits in length. This address is resolved as a displacement (aaa) from the base register (b), which contains a 24-bit address. This defines a 24-bit address with only 16 bits in the instruction. Also, because all addresses are relative to a base register set by the program, the object code is relocatable in its executable format. Thus, less overhead is involved during the program load, as only address constants must be resolved at this time.

However, in the S/360, once a program is loaded it may not be moved. The loader resolves address constants with absolute addresses and the program execution places absolute addresses in the required base registers. From that point on, the program is fixed in the real memory locations where it was loaded. Thus, for swapping to work, a program would have to be swapped back into these same storage locations.

IBM developed two multiprogramming systems on the S/360 series. The first, MFT, divided memory into a number of fixed partitions. The second, MVT, dynamically assigned real memory to jobs as they were started.

Neither of these system designs was flexible enough for a time-sharing environment. In MFT, the fixed partitions made swapping a possibility but required a preallocation of memory, which severely constrained the batch workload. In MVT, swapping was not practical because a job had to be swapped back into exactly the same real memory it was first loaded into, and that memory had to be contiguous.

The first IBM time-sharing system combined ideas from each of these systems. The first Time-Sharing Option (TSO), released with MVT in the early 1970s, allocated a number of fixed time-sharing partitions, with the remainder of real storage dedicated to batch.

Figure 2-21 shows a memory allocation for an MVT system. The System Control Program (SCP) occupied the high and low addresses, with the area in between dynamically assigned to jobs.

In MVT, the memory between the executing programs was available for execution. Eventually, 50 percent of the available memory is lost to the type of memory fragmentation shown in Figure 2-21. This type of fragmentation is called EXTERNAL FRAGMENTATION. In MVT and in other similar systems, the real memory lost to external fragmentation was shown to average around 33 percent. This represented a serious constraint on throughput, as well as requiring a substantial effort for the installation in trying to minimize the occurrence and impact of this fragmentation.




Figure 2-21. An MVT Real Memory Allocation

DYNAMIC ADDRESS RESOLUTION

This discussion of dynamic addressing outlines the benefits of using a dynamic scheme in implementing a virtual system. The mechanics of dynamic addressing in MVS is discussed in Section 2.2.4.1.2. Address resolution performed as each instruction is executed solves the problem of code and data relocation. Two schemes were implemented: segmentation systems and paging systems. A segmentation system divides each program into a number of parts called segments. Segments of a program may be relocated in real memory. The number of segments a program could be divided into was small. On the GE635, for example, there were seven segmentation registers. As this scheme still had the external fragmentation problem illustrated by Figure 2-21, it never gained wide use.

In a paging system, all real memory is divided into equal parts (pages) of a fixed length. Each active program in virtual storage is "paged in" to real storage in order to gain access to the CPU. A paging system allows the virtual pages making up a job to be placed anywhere in real memory. In a paging system, there can be no external fragmentation, because each program in memory occupies an integral number of pages. Each program has some wasted space, because allocation is by pages. This is called internal fragmentation. This problem is minimized for a given system by the selection of a page size and by the techniques used for dynamic storage allocation within a program. Early work in this area with IBM virtual systems is described in the IBM Journal of Research and Development.

One of the goals of a virtual system is to relieve application programs of memory management concerns. In non-dynamic systems such as MVT, a great deal of system design went into making a program use memory more efficiently through the use of overlay structures and dynamic loading facilities. In a virtual memory system, a program is given a large amount of virtual memory to use. The operating system then allocates real storage to jobs as required. Figure 2-22 shows a simplified view of the virtual storage allocation in an MVS/370 (pre-MVS/XA) system running four programs. To each program, it appears as if the entire virtual storage of 16 megabytes is available to it (less the areas used by the control program).




Figure 2-22. MVS/370 Virtual Memory Allocation

An implication of the MVS/370 24-bit addressing scheme is that the size of a user application program is constrained by the limitations of 16 megabytes minus the size of the control program. As user application programs grew larger and larger, it became evident that an addressing scheme that allowed only 16 megabytes of virtual storage addressability was inadequate.

To provide the necessary virtual storage constraint relief, allowing users to write applications that are substantially larger than those possible in MVS/370 systems, MVS/XA (extended addressing) was developed. By expanding the size of the address to 31 bits, jobs can address up to 2 billion bytes (2 gigabytes) of virtual storage. To maintain compatibility with MVS/370 and yet provide the expanded addressability, MVS/XA provides regions of virtual storage that are defined below the 16 megabyte line and "extended" portions of these regions are defined above the 16MB line.

Figure 2-22 shows conceptually how this is done. MVS/XA treats both the portion below the 16MB line and the extended portion above the 16MB line as one logical area. As with the example in Figure 2-22, each job running in an MVS/XA system would have its own copy of the Private and Extended Private areas.




Figure 2-23. MVS/XA Virtual Memory Allocation

You can determine whether the program, or modules, use 24 bit or 31 bit addressing via the AMODE attribute which is specified to the assembler or the linkage editor. The default if the AMODE attribute is not specified is 24 bit addressing.

You can also determine whether the modules are to be loaded into the User Region (below the 16MB line) or the Extended User Region (above the 16MB line) via the RMODE attribute for the assembler or linkage editor. The default if the RMODE attribute is not specified is to load the module into the User Region below the 16MB line. This allows programs written for MVS/370 to run in MVS/XA with generally no modification.

Figure 2-24 shows how real memory might be allocated to the four jobs concurrently executing in Figure 2-22. First, no real memory is lost to external fragmentation. If another job starts and no real memory is available, either some pages may be taken from other jobs to make real memory available, or one or more of the active jobs may be removed from memory (in MVS, swapped out).

Real memory is more efficiently used because idle portions of the program (either code or data) are not kept in real memory. In Figure 2-24, the page numbering indicates the page number in the program's virtual storage. Thus, in MVS, where the page size is 4096 bytes, page 1 is the first 4K of the program, page 2 the second 4K, and page n is the nth 4K section of the program.

We see from Figure 2-24 that only referenced parts of the program need to be resident and that these need not be contiguous in real memory. In theory, a program could be executing (efficiently) with only one page in memory at a given time.




Figure 2-24. An MVS Real Memory Allocation

2.2.4.1.2 Paging Concepts

Having settled on paging as the method of providing virtual storage, three issues that have a large impact on the final design must be considered. These are page size, the method of mapping a virtual to a real address, and the determination of which pages should remain in real storage.

PAGE SIZE DETERMINATION

In MVS, to remain compatible with S/360 and S/370 storage protection features and existing storage allocation software, a page size that was a multiple of 2 KB was dictated. MVS chose a page size of 4 KB, accepting that potentially 4095 bytes of storage could be wasted.

The choice of page size is a compromise between the amount of storage lost to internal fragmentation and the size of the tables that translate virtual addresses to real addresses (the Page Tables).

Another issue in considering the page size is the page transport time, that is, the time required to write the page to auxiliary storage. With the 16 MB virtual address space provided by MVS through SP1.3, the 4 KB page size has worked well. In later systems, even though the amount of virtual storage has been increased to 2 GB, the page size has remained 4 KB. This size remains in use today with z/OS.

The remainder of this chapter uses the terms page, frame, and slot in a very specific manner as defined below:

- o Page - A unit of virtual storage (4 KB or 4096 bytes in MVS).
- o Frame - A 4 KB unit of real storage that contains data for some virtual address (a page).
- o Slot - A 4 KB unit of space (on a direct access storage device) that holds the data contents of a page when it is not occupying a frame in real storage.

VIRTUAL ADDRESS TRANSLATION

The second issue to be considered is the mapping of virtual to real storage addresses. A simple structure would be to have each virtual address be a page number and displacement within the page. The problem with this implementation is

that it requires the entire page table to be allocated at all times. In MVS/370, the page table would have been 4096 pages for each active task; in later operating systems the number would have become 128 times greater.

The solution to this problem is to divide the virtual address into fixed segments. In MVS/370, there are 256 segments, each representing 64 KB of the virtual address space. Each segment then is defined by 16 pages. In MVS/XA and later systems, to accommodate the much larger virtual storage, there are 2048 segments, each representing 1 MB of address space. Each segment is divided into 256 pages.

In all systems, the virtual address is a triplet (s,p,d), that is, a segment number, page within the segment, and displacement within the page. The control program then only needs to have page tables for active segments. The page tables define all valid pages belonging to the program. If the page is in real memory, the page table contains the starting address of the page. If the page is not currently in real storage, the page table contains a status bit that indicates this fact. A complete description of the hardware supporting virtual address translation is found in the Principles of Operation.

Figure 2-36 shows how the segment and page tables map virtual addresses to real addresses. In MVS, if we do not consider the extended virtual addressing schemes, a virtual address is 24 bits, or 6 hexadecimal digits of the form:

sspddd

where: ss = segment number (0 - 255)
p = page number (0 - 15)
ddd = page offset (0 - 4095)

If the 31-bit addressing scheme is used, a virtual storage address is of the form:

ssspddd

where: sss = segment number (0 - 2047)
pp = page number (0 - 255)
ddd = page offset (0 - 4095)

Note that the leftmost bit of the sss field is not used.

Starting with MVS/SP 1.3, support was added for real addresses greater than 24 bits in length. This is

accomplished by adding two bits to the real addresses in the page tables. The changed support thus provides for 64 MB of real storage. When frames are assigned to pageable storage areas, an available frame from either below or above 16 MB (the maximum 24-bit address) can be used. However, when a page is fixed (usually for I/O to be done), if the page is currently occupying a frame above the 16 MB address, it must be moved to a frame below 16 MB, because the channels only support 24-bit addressing.




Figure 2-25. MVS/370 Virtual to Real Address Translation

PAGE MANAGEMENT STRATEGIES

The final issue is to determine which pages should be in real storage. There are two parts to this question: how should pages be brought into real memory, and which pages should be replaced when more real memory is required? These algorithms

are called fetch and replacement strategies.

Several fetch strategies have been used in various systems. Some of these are:

- o To load all pages of a program before execution.
- o To try to predict which pages will be used and to fetch these pages in advance of their reference.
- o To bring in pages as they are needed.

In MVS, a combination of these ideas is used. When a program is initiated, all of its pages are brought into real memory by the program load facility (fetch). This process works in the following manner:

- o Virtual storage is allocated for the program with the MVS virtual storage management component (VSM), that is, GETMAIN.
- o Assignment of real storage frames to back up the virtual storage is done by the real storage management component (RSM).
- o The virtual storage (that is, the range of virtual addresses to be occupied by the program) is page-fixed to prevent page stealing during the loading process.
- o The program is now brought into real storage from the load library (a direct access data set such as SYS1.LINKLIB) into its assigned virtual locations.
- o After the I/O to read the program into virtual storage is completed, the pages are freed, that is, the normal demand paging process is now allowed to operate on the program just loaded.

At this point, the pages are subject to the normal page replacement process. However, under normal circumstances, some time must pass before this will occur, since all the pages are "recently referenced" at this point.

Once execution begins, pages that have been paged out are brought into real memory as they are referenced. This scheme is called demand paging, that is, pages are brought into memory as they are "demanded."

The idea of bringing in the entire program initially is to

avoid forcing the program to establish its initial set of necessary pages by repeatedly demand paging until it can get down to work. It also serves to provide a means of getting the program backed up in auxiliary storage through the existing page replacement facilities, rather than directly allocating Auxiliary Storage Manager (ASM) slots and writing the program to them.

As part of the MVS page replacement process, a queue of free frames is maintained by RSM. This queue, called the Available Frame Queue (AFQ), serves several functions:

- o It reduces the time required to resolve a page fault, since page fault resolution does not (normally) have to wait on a page-out to complete.
- o The length of the AFQ determines how often page replacement processing is done.
- o By maintaining a reservoir of free frames, MVS can more easily meet sudden demands for relatively large numbers of pages (such as that of the program load process just described and for the swap-in process.)

When all the frames comprising real memory are allocated and a program requires a page that is not currently in memory, a page must be selected to be written back to auxiliary storage to make room for the new page. The goal of the page replacement process is to select a page for replacement that will not be referenced again for a relatively long time. Historically, several methods have been tried:

- o First In, First Out - Select the oldest page in the system for replacement.
- o Working Set - Only replace pages that a program has not used for a given period of time.
- o Least Recently Used - Replace the page that has not been referenced for the longest period of time.

Modified pages that are selected for replacement must be written to the auxiliary storage to preserve the changes. If a page that is queued to be written is referenced again before the actual I/O occurs, the page is said to be reclaimed. A high reclaim rate indicates a problem, since this means the page replacement strategy is not working properly (that is, it is selecting pages to write out that are being quickly and/or frequently rereferenced).

MVS uses a global Least Recently Used (LRU) algorithm. Theory and practice have shown that this method also gives a good approximation to the ideal: to replace a page that is not likely to be referenced again soon. It turns out that the application of this strategy also gives a practical definition of a program's "working set."

The concept of a working set was created to explain a particular problem. Early virtual systems were subject to a phenomenon called "thrashing," which occurred when the page replacement strategy broke down. This happens when replaced pages are immediately referenced (and therefore must immediately be brought back into real storage).

Peter Denning explained thrashing in terms of a program's working set. That is, at any point in time, each program has a minimal number of pages that it requires to run in a virtual environment. When the program's working set is available, all is well. According to the theory, when pages are stolen from a program's working set, paging increases exponentially until the system is doing nothing else. Thus, thrashing is characterized by an excessive amount of CPU time being expended by the RSM and ASM components of the operating system.

The MVS real storage management algorithms do a good job of preventing thrashing from occurring. When it does occur in MVS, it is usually because the system minimum multiprogramming level (MPL) as specified in the Installation Performance Specifications is too large. Working set concepts are discussed further in Section 2.2.4.2.2. The MVS storage management algorithms and storage controls are discussed in Section 2.2.4.3.

2.2.4.2 MVS Virtual Storage Overview

This section describes the individual parts of the MVS virtual memory, identifying those parts of the virtual memory that are permanently page-fixed (that is, cannot be paged out). The size of these areas affects your overall real storage management strategy, because the total size of these areas is subtracted from the page frames available to the workload. It also affects the size of the private area (the amount of virtual memory available to programs).

MVS/370 provided 16 MB of virtual storage to each address space. This storage is divided into the three areas: the common area, the private area, and the nucleus area.

In MVS/XA and later systems, 2 GB of virtual storage are available to each address space. This storage is divided into four areas: the common area below the 16 MB line, the private area below the 16 MB line, the extended common area above the 16 MB line, and the extended private area above the 16 MB line. In current systems the nucleus, which spans the 16 MB line, is conceptually part of the common/extended common area.

In considering the virtual address space, there are both real storage and virtual storage issues to be considered. This section defines each of the parts of the virtual address space, and details the real and virtual storage issues pertinent to each part.

Real storage issues arise because portions of the virtual address space remain fixed in real storage at all times. A page fixed in real storage is one that remains in real storage and will not be paged out. Using Figure 2-27 as a reference, the fixed portions of real storage in MVS/370 are:

- o The entire nucleus area.
- o The system queue area (SQA) and parts of the common storage area (CSA) from the common area.
- o The local system queue area (LSQA) from the private areas of resident address spaces.

Using Figure 2-28 as a reference, the fixed portions of real storage are:

- o The nucleus and extended nucleus area.

- o The LSQA and extended LSQA from the private areas of resident address spaces.
- o The fixed link pack area (FLPA) and extended FLPA (in MVS/370 the FLPA was an extension of the nucleus).
- o The SQA and extended SQA.

Note that in both MVS/370 and current systems, LSQA is swappable. Thus, the real storage occupied by LSQA and the private area fixed pages are also determined by the number of non-swappable tasks and the number of active tasks (the system MPL).

The size of the nucleus and common areas must all be considered, because the total virtual space is limited. Starting with the Systems Product releases of MVS, the size of the available private area is the remainder after the nucleus and the common area are subtracted.

In earlier systems, the size of the common area was fixed at 8 MB. This could be changed by modifying a constant in one of the IPL modules. In the larger MVS/370 systems today, private areas are as small as 5 MB. In these systems, reduction of the virtual sizes is an important issue.

MVS itself requires parts of both real and virtual storage. In determining how much storage is actually available to users, the basic requirements of MVS must be subtracted from the total configuration. The real storage manager (RSM) controls how real storage is allocated to system and user address spaces that are swapped in for execution.

The primary MVS requirements for real storage are:

- o Fixed storage for resident supervisor code and basic system control blocks. This includes, for example, the code for the System Resource Manager (SRM); the real, virtual, and auxiliary storage management routines (RSM, VSM, ASM); the dispatcher; basic systems services; and control blocks for defining the real/virtual storage mappings and for address space control.
- o Fixed storage requirements to support user address spaces. This includes the page and segment tables for the address spaces, address space control blocks, and address space-related SRM/RSM/VSM/ASM control blocks.

Figures 2-27 and 2-28 show the virtual storage layout for all

MVS-based systems. The System Storage Usage Report described in Section 2.2.3.2 provides the capability to track most of the different parts of virtual memory. In storage-constrained systems, a great deal of real storage can be made available by tailoring the system's use of real memory.

Section 2.2.4.2.1 describes each of the components of the MVS virtual memory. As shown in Figures 2-27 and 2-28, there is a private area (and an extended private area in later systems) for each address space, that is, one for the Master Scheduler, JES2/3, VTAM or TCAM, any started system task, each batch job, and each TSO user. Each address space in memory has fixed storage requirements: its LSQA and address space related control blocks in SQA. When planning for the amount of pageable storage, the fixed storage required to support swapped-in address spaces must be considered.

Part of real storage is partitioned into a system preferred area to support storage reconfiguration for multiprocessors. Allocations for SQA, LSQA, and fixed page assignments (including frames for V=R jobs) are made from the system preferred area if possible, in order to facilitate storage reconfiguration.




Figure 2-27. MVS/370 Virtual Storage Layout

2Gb +---		=====	=====	=====
Extended Private Area <		EXTENDED LSQA	EXT LSQA	EXT LSQA
		-----	-----	-----
		EXTENDED CSA	EXT CSA	EXT CSA
		-----	-----	-----
		EXTENDED 229/230	EXT 229/230	EXT 229/230
		=====	=====	=====
		MASTER	addr. space	addr. space
		SCHEDULER	1	2
		PRIVATE AREA	extended	extended
		EXTENDED	user region	user region
		USER REGION		
+---		=====	=====	=====
Extended Common Area <		EXTENDED CSA		

		EXTENDED PLPA/ FLPA/MLPA		

		EXTENDED SQA		
+---		=====		
16 MB +---		=====		
Common Area <		NUCLEUS		

		SQA		

		PLPA/FLPA/MLPA		
Private Area <		-----		
		CSA		

		LSQA	LSQA	LSQA
		-----	-----	-----
		SWA	SWA	SWA
		-----	-----	-----
		229/230	229/230	229/230
		=====	=====	=====
		MASTER	address	address
20K		SCHEDULER	space	space
		PRIVATE	1	2
		AREA	user region	user region
		USER REGION	-----	-----
			Sys. region	Sys. region
4K +---		=====	=====	=====
Common <		PSA		
		=====		
0K +---		=====		

Figure 2-28. Current systems Virtual Storage Layout

The following sections explain the portions of the virtual storage defined in Figures 2-27 and 2-28, describe the MVS real storage algorithms, and describe a method of determining the working set of a program:

- 1 - MVS Virtual Storage Layout
- 2 - MVS Working Set Determination

2.2.4.2.1 MVS Virtual Storage Layout

This section defines the different portions of MVS virtual storage. (The basic reference for this material is the Initialization and Tuning Guide.) Real storage management considerations specific to each area are included following the description of each section.

SYSTEM QUEUE AREA (SQA and Extended SQA)

Generally speaking, the system queue area (SQA) is an area containing global system control blocks. At system initialization, a minimum of 64 pages are allocated to SQA. A number of additional segments may be added by PARMLIB or SYSGEN options. These parameters specify the upper limit on the number of frames that will be allocated to SQA. Real frames are allocated as required. SQA is fixed storage acquired and freed by various MVS components (for example, SRM, RSM, VSM, ASM, ENQ/DEQ). If the system requirement for SQA exceeds the amount of pages specified by the SQA parameter, the system attempts to allocate virtual SQA from the common service area (CSA).

In MVS/370, if the amount remaining for allocation purposes falls below 6 pages, the creation of new address spaces is suspended. In later systems the threshold is 8 pages below 16 MB.

You should monitor the size of SQA on a long-term basis for memory planning considerations. System upgrades, maintenance, and new program products may all increase the size of SQA. You should also review the SQA on a daily basis to ensure there are no SQA growth problems.

SQA pages are allocated from the system preferred area if possible. If preferred area pages are not available, and if all attempts to make a page frame available (both by page stealing and page relocation) have failed, then the reconfigurable storage (storage available to be assigned to a processor that is varied offline in a multiprocessor) is reduced by one storage unit (a processor-dependent value, 4 MB on a 3081) and the storage unit is marked as being part of the preferred area. As a last resort, a page is allocated from the V=R area.

PAGEABLE LINK PACK AREA (PLPA and Extended PLPA)

The Pageable Link Pack Area contains the code for SVC (supervisor call interruption) routines, access methods, and

other system routines. A packing algorithm is used during PLPA initialization to minimize fragmentation of this area. In addition, a packing list should be provided to minimize the working set for PLPA by placing high activity modules, or modules that are used together, in the same page.

Expected sizes of PLPA range from 2 to 4 MB. In systems with virtual space sizing problems, the removal of unused modules from PLPA offers the potential for large gains. A field-developed program (FDP) that provides LPA analysis, MVS LPA Optimizer, SB21-3011, is available to help with this process.

MODIFIED LINK PACK AREA (MLPA and Extended MLPA)

The modified link pack area can be used to temporarily include reentrant modules from the linklist libraries in the PLPA. The intent is to provide a testing capability. The MLPA is serially searched after the FLPA list, but prior to the LPA directory search. The LPA directory searches for PLPA modules with a hash algorithm. Thus, a long and "permanent" MLPA list negates the benefits of the more efficient PLPA lookup.

A large MLPA list also causes excessive dispatching overhead because each MLPA search/update requires the CMS lock. A long MLPA list also increases SQA size, because a Contents Directory Entry (CDE) and Extent List (XTLST) are required for each entry. The list of modules to be included in MLPA is in a member of SYS1.PARMLIB, IEALPAnn. Modules may be included from any linklist library. When modules are included from libraries other than SYS1.LPALIB, SYS1.SVCLIB, or SYS1.LINKLIB, use SYS1.LINKLIB as the data set name for the module.

FIXED BLDL LIST (BLDLF)/PAGEABLE BLDL LIST (BLDL)

A BLDL list is an in-storage copy of all, or a portion of, the directory for a linklist data set. The fixed BLDL list should point to modules that are frequently used but cannot be placed in PLPA because they are not reentrant or because their size and/or frequency of use do not justify placement in PLPA or MLPA. Either a BLDLF or BLDL list may be defined in MVS/370. In later systems, the BLDL/BLDLF lists have been eliminated.

The BLDL list should also include modules that are referenced by few address spaces and modules that are infrequently used by many address spaces. The primary purpose of the BLDL list

is to reduce the I/O overhead from linklist directory lookups. There is no advantage in specifying a pageable BLDL list. The BLDL list is searched after the resident lists (FLPA, MLPA, and PLPA), but before the linklist library directories are searched.

Establishing the content of a BLDL list is a time-consuming task that does not address shifts in the character or volume of the system's workload. For example, a BLDL list designed to optimize TSO work is not of much use at night when the production batch is run. A better solution to this problem may be to use one of the dynamic BLDL packages that change the content of the list based on frequency of module use.

The size of a BLDL list is not an issue as it requires very little storage.

SYSGEN PREFIX STORAGE AREA (PSA)

The prefix storage area (PSA) represents the first 4K of storage starting at address 0. There is one PSA for each processor on a multiprocessor (MP) or dyadic processor (3081). The PSA generally contains CPU-dependent, hardware-related information. For systems with multiple processors, an uninitialized copy of the PSA is allocated to initialize the PSA for each CPU brought online after system initialization (IPL). There is no extended PSA allocated above the 16 megabyte line.

COMMON SERVICE AREA (CSA and Extended CSA)

The Common Service Area (CSA) contains pageable and fixed data areas that must be addressed by more than one address space. The determination of whether a CSA frame is pageable or fixed is based on which virtual storage subpool the area is to be allocated from. Subsystems such as TCAM, VTAM, JES3, HSM, and IMS use large amounts of CSA to pass data from one address space to another. The virtual size of CSA is a consideration if there are problems with the size of the private area.

LOCAL SYSTEM QUEUE AREA (LSQA and Extended LSQA)

The Local System Queue Area (LSQA) is allocated from the private area and is fixed in storage while the address space is resident. It contains local system control blocks. LSQA is allocated from the system-preferred area following the same scheme described for SQA allocation. The LSQA for each address space is unique. The sum of LSQA allocations for

address spaces that are resident represents a block of fixed pages that is not available for demand paging. There are two ways to minimize the amount of storage used by the LSQAs in a storage-constrained system. First, reduce the system MPL. Second, limit the number of non-swappable address spaces.

SCHEDULER WORK AREA (SWA and Extended SWA)

The scheduler work area (SWA) is a pageable area allocated at the top of the private area that contains control blocks that exist over the life of the address space. The SWA eliminates the job queue that existed as a direct access data set in earlier IBM operating systems.

There are no real storage considerations associated with the SWA. The only virtual storage consideration is that this area is in the user private areas.

SUBPOOLS 229/230 (and Extended 229/230)

The subpools 229/230 area contains pageable space for system control blocks within a virtual address space. These subpools are used to provide data areas that may be referenced by MVS components with an appropriate storage protect key.

The MVS/SP Tuning Cookbook recommends that LSQA allocations that do not need to be fixed be made from these areas instead. For systems that are severely memory-constrained, this may provide some temporary relief. Because system code modification is required, and this change may impact installation exits that may reference the control blocks, this should be evaluated against the other ways to make more real storage available.

SYSTEM REGION

The System Region is reserved for use by the address space management functions of MVS. It comprises the bottom 16 KB of each private address space other than the Master Scheduler address space. In the Master Scheduler address space, the system region is a maximum of 200 KB. In later systems there is no extended version of this area.

PRIVATE USER REGION (Private and Extended Private)

The Private User Region contains user code and data. In MVS/370 systems, its virtual size ranges from 5 (or less) to 12 MB, depending on the size of the other areas

comprising the virtual space. In current systems, with a virtual address space of 2 GB, its size is considerably larger than that found in MVS/370.

The minimum acceptable size for the private area is determined by installation needs. For systems that are having problems with the size of the private area, scaling down the ASM configuration, reducing the virtual size of PLPA, and minimizing the number of concurrent address spaces is required.

FIXED LINK PACK AREA (FLPA and Extended FLPA)

The Fixed Link Pack Area (FLPA) is intended for modules that execute more effectively when page-fixed. In MVS/370 systems, the fixed link pack area is an extension of the nucleus. Its size is included in the nucleus figure reported by the Real Storage Usage Report for MVS/SP 1.3 systems. In current systems, it is allocated separately and is no longer an extension of the nucleus.

For storage-constrained systems, only essential modules should be kept in this area. Because the paging algorithms normally keep frequently used modules in real storage, this area should be used only to hold modules that are not frequently used but are needed for fast response to some terminal-oriented action.

Modules are loaded into the FLPA area in the order specified by the IEAFIXnn member of SYS1.PARMLIB. Thus it is important to minimize the amount of space lost to fragmentation. One method of doing this is to order the fix list by descending module size.

NUCLEUS and Extended Nucleus

The nucleus contains the nucleus load module and some key control blocks. The IPL extensions are discussed separately, so this discussion is limited to the contents of the base nucleus.

From the view of recovering real storage from the nucleus, the I/O generation should be reviewed periodically to remove unit control blocks (UCBs) for devices that are no longer used. Because the nucleus high address is rounded to a segment boundary (64K) in MVS/370 systems, a larger than expected amount of storage may be recovered. On the other hand, in MVS/370 systems, you may find that additional modules may be added to the FLPA list, or the BLDL list may

be extended "for free" due to the already existing dead space that results from this rounding.

In later systems, because the BLDL list does not exist and because the FLPA is allocated separately, this possibility for using the wasted space does not exist.

After the installation of upgrades or system maintenance, you should review the impact on the size of the nucleus and all fixed areas. For example, after the addition of more real storage, the size of the nucleus increases, because the page frame tables are in the nucleus area. Addition of 8 MB, for example, would increase the nucleus by 32 KB (2048 frames and 16 bytes per entry in the frame table).

2.2.4.2.2 MVS Working Set Determination

MVS/SE2 introduced a facility to guarantee a performance group a fixed amount of real memory. To use this facility, storage isolation (fencing) it is desirable to be able to characterize the working set of a program that is to be fenced.

MVS maintains a measure of a program's working set, but to be able to use fencing fully, you would like to be able to estimate the paging that would be expected with various amounts of real storage available to the program.

With this information, you can use fencing to minimize the page faulting of a program, or you can allocate the amount of real storage expected to yield a desired paging rate. The Real Storage Analysis portion of the CA MICS Performance Manager Option provides reports and plots that support the working set characterization method described in this section.

You could characterize the working set requirements of a program by repeatedly executing the program, limiting to it different amounts of real storage in fixed increments. Starting at 50 pages and adding 50 pages with each execution, you should expect (in theory) to see a curve something like that shown in Figure 2-29.

Note that this procedure gives an estimate of the theoretical working set of the program. In applying storage isolation, an estimate of the theoretical working set is what is desired. When real storage is in plentiful supply, programs that are storage isolated are allowed to accumulate more pages, just as the other address spaces are. In times of increased storage demand, however, it serves no purpose to allow the storage- isolated tasks to retain more storage than they require, as they will, in most cases, reach the same paging rates after some period of time.




Figure 2-29. A Working Set Characterization

Arthur Petrella and Harold Farrey proposed a functional definition of the working set in MVS that has worked well in practice. The working set determination method used here is based on their ideas.

The working set is, at any point in time, the number of pages referenced during the last interval (t). Since the real memory available to the program differs over the course of execution, as well as from run to run, a function describing the working set must represent the dynamics of the available real storage as well as the characteristics of the program. Thus, a probabilistic representation is required. Petrella and Farrey proposed a binomial probability function.

A set of events may be represented by a binomial distribution if three conditions are met:

- o There are only two things that can happen (success or failure).
- o The probability of success (p) is constant from trial to trial.
- o There are n independent trials. That is, the outcome of one trial does not depend on the outcome of any of the others.

If you define a trial as the occurrence or non-occurrence of a page fault, then conditions 1 and 3 are certainly met. And if a measurement interval of length t is chosen such that there can be at most one page fault in each interval, then condition 2 is also met.

Now, given that the probability of success (a page fault) is small, if you consider the probability of N page faults over a time interval T , with T much larger than t , the binomial distribution can now be approximated by the Poisson distribution shown below. For a more complete discussion of the Poisson process, refer to page 117 of Probability, Statistics, and Queueing Theory by Arnold O. Allen.

$$P(T) = e^{-k} \frac{k^N}{N!}$$

So, we expect that Figure 2-29 is an exponential relation. With this relation, we can estimate the results of the fencing experiment from the standard SMF data collected during regular production runs of the program. In the fencing experiment, over any single run what we are really measuring is the expected value of this function. CA MICS provides three measures that can be used to plot the relation in Figure 2-29:

- o PGMPGSEC - in the batch program file is the product of the number of pages used by the program and CPU time.
- o PGMPGIN - is the number of page-ins during the program's execution.
- o PGMPRCLM - is the number of page reclaims during the program's execution.

From these variables, CA MICS computes average working set size in pages (PGMAVWSS), and the Real Storage Analysis component calculates the paging rate as PGMPGIN + PGMPRCLM. Thus, by retrieving and analyzing several executions of a program, the relationship suggested by Figure 2-29 can be reproduced.

Because we are dealing with a relatively small number of samples, it is important to be able to test the hypothesis, "Does the data have an exponential relation?" and also to be able to project an optimal amount of memory to assign to the program. An exponential function may be transformed into a linear one by:

$$\ln e^{(mx+b)} = mx+b$$

We add one to the CPU page fault rate to allow for the case where there is no paging. Now, the relation shown below is linear.

$\ln(1 + \text{page faults/CPU second})$ vs. working set size

If these values are plotted for a set of executions of a program, the linear relation shown above can be estimated by regression techniques. The X intercept of this plot approximates the virtual memory used by the program. An even more interesting result is that the slope of the line indicates whether a program that has already been fenced has more real storage than is required for its working set via the storage isolation definitions.

The final step in the process is to investigate the quality of the linear model. Does the data indeed have a linear relation? How well does the regression model describe the variation in the data? To do this, the SAS procedure REG can be used.

2.2.4.3 MVS Real Storage Control Mechanisms

When looking at the real and virtual storage mechanisms of MVS, it is important to remember that MVS is not a paging system. That is, MVS does not manage real storage by paging; real storage management is done primarily by swapping. Paging in MVS provides several critical improvements over the earlier IBM systems:

- o More efficient use of real storage. Paging completely eliminates the external fragmentation of the MFT and MVT systems.
- o Provides the application with a large virtual memory. MVS largely removed issues of memory management from applications.
- o Improvements in data and system integrity. A dynamic addressing scheme makes this much easier to accomplish.
- o Increase in the number of concurrent jobs. The S/360 and S/370 architecture has a maximum of 16 storage protection keys. In the MFT and MVT systems, this would be a severe limitation on batch throughput. In MVS, because of the protection afforded by dynamic addressing, all user (problem state) work except for V=R jobs may use the same protect key.
- o Supports Swapping - Swapping in a non-paging environment would be much more complex because of address relocation issues.

Because of this design, the MVS paging algorithms are vastly different (and more complex) than those of a system whose primary means of real storage control is through paging. Management of storage is divided among the three subsystems (RSM, ASM, and VSM) described below:

REAL STORAGE MANAGER (RSM)

The Real Storage Manager (RSM) maintains the control blocks that translate real and virtual storage address, makes the determination of working set contents when an address space is swapped, and performs the page replacement operations.

AUXILIARY STORAGE MANAGER (ASM)

The Auxiliary Storage Manager (ASM) manages the page and swap data sets that comprise the virtual space. The page data

sets provide the auxiliary storage that backs up the storage frames. Prior to the introduction of the extended swap facility (provided either by the Installed User Program (IUP) or with MVS/SP 1.3), the swap data sets contained only the LSQA frames from swapped-out address spaces. The private frames were paged out. With the extended swap processing, the LSQA pages and the working set of the address space are written out together to the swap data sets. This concept is discussed in detail in Section 2.3.

VIRTUAL STORAGE MANAGER (VSM)

The Virtual Storage Manager (VSM) provides a map of the private area virtual storage within each address space. VSM interfaces with RSM to handle the dynamic allocation of virtual memory.

STORAGE MANAGEMENT TECHNOLOGY

MVS uses a least-recently-used (LRU) page replacement algorithm to keep a number of free pages to be used for page allocations and swap-ins. Through the progression of the MVS Systems Extensions and System Product Releases, the real storage algorithms have been refined and additional capabilities have been added. The remainder of this section defines the operation of the page replace and storage control algorithms and traces some of the historical development of these processes.

- 1 - Swapping
- 2 - UIC Update Process
- 3 - Page Steal Algorithm

2.2.4.3.1 Swapping

The resource use functions of the SRM attempt to optimize the use of system resources on a systemwide basis. It does this through multiprogramming level adjusting, CPU load balancing, I/O load balancing, storage load balancing, and demand or logical swapping.

Swapping is the primary method used by the SRM to effect this control. Using IPS and OPT parameters and system status information, the SRM determines which address spaces should have access to system resources.

There are a number of reasons for swapping. The first three described below control domains and the address spaces within a domain, while the next four control system performance and throughput.

Unilateral Swap In:	If the number of active address spaces for the domain is less than the value specified in the IPS or less than the SRM deems optimal, the SRM swaps in address spaces for the domain.
Unilateral Swap-Out:	If the number of active address spaces for the domain is greater than the value specified in the IPS or greater than the SRM deems optimal, the SRM swaps out address spaces for the domain.
Exchange Swap:	All address spaces within a domain compete with one another for system resources. When an address space has consumed its allotted resources and another address space in the same domain is waiting to be swapped in, the SRM swaps out the one in real storage and swaps in the other.
Swaps Due to Storage Shortages:	If auxiliary storage or pageable frame shortages occur, the SRM swaps out the address space that is acquiring auxiliary storage at the fastest rate or the address space with the greatest number of fixed frames. It continues this operation until the shortage is corrected.
Swaps Due to Wait States:	In certain cases the address space signals the SRM that it is a candidate for swap-out. This occurs when it can be determined that the address space will not become eligible for execution for a relatively long time. Examples are an ENQ for a resource currently held by a

swapped-out address space or an STIMER wait of greater than 0.5 seconds.

Request	The system may request that an address space be
Swap:	swapped out. For example, the VARY STOR,OFFLINE command requests the swap-out of address spaces that occupy frames on the storage unit to be taken offline.
Transition	A swap occurs when the status of an address
Swap:	space changes from swappable to non-swappable. This special swap prevents the job step from improperly using reconfigurable storage.

The SRM adjusts the multiprogramming level by adjusting the number of address spaces that are held in real storage for each domain. When the paging rate becomes excessive, the SRM reduces it by swapping out address spaces to reduce paging. If the rate is too low, the SRM swaps in address spaces to increase it.

The SRM adjusts the CPU load level by swapping in eligible address spaces if the CPU utilization is too low or swapping out address spaces if the CPU utilization is too high.

The SRM adjusts the I/O load level by swapping in or out address spaces to maintain the proper utilization of the system's logical I/O channels.

The SRM adjusts the storage load level by swapping in or out address spaces to keep real memory properly utilized.

The SRM effectively utilizes real storage and reduces channel and processor overhead by logically swapping some address spaces. Rather than swapping them to auxiliary storage, they are left in real storage whenever possible. This is done if the system determines that the address space will be ready to execute in a relatively short period of time.

With the advent of RMF 4.2, four new swap reasons are reported, which are used to manage work flow and resource commitment. The SRM monitors the utilization levels of various system resources and attempts to bring these levels to within site-specified limits by swapping address spaces out of central storage. In addition, the SRM may respond to an address space experiencing a wait condition during an APPC conversation by swapping the address space.

Swapping may be performed in order to improve the system

paging rate. If the SRM has been set up to manage the system paging rate, and that rate rises above the maximum acceptable level, the SRM will cause address spaces to be swapped out in an attempt to lower the rate.

Swapping may also be performed in order to improve central storage usage. If the SRM has been set up to manage central storage, and demand for storage rises above the maximum acceptable level, the SRM will cause address spaces to be swapped out in an attempt to lower the demand.

Swapping may occur to make room in central storage to swap in an address space that has been swapped out too long. The SRM monitors the duration of swap-outs, and if no other condition has made it possible to swap in an address space that has been swapped out for a long time, the SRM will make room in central storage for it by swapping out other address spaces.

During an APPC conversation, an address space issuing an LU6.2 call to the APPC address space can experience a wait condition. When this occurs, the SRM can then make a decision whether to swap the calling address space out or not.

2.2.4.3.2 UIC Update Process

Maintaining a record of the time since each page was last referenced is central to the implementation of a LRU page replacement algorithm. This function, called the UIC (unreferenced interval count) update process, is actually implemented by the same routine that performs the page stealing function. We will discuss the two functions separately in order to make it easier for you to follow the developmental changes to each of the processes.

The UIC records the number of seconds since the page has been referenced. The system UIC is thus the age in seconds of the oldest page in storage. In Figure 2-30, the rates for UIC updating for various values of the system UIC are presented for both swappable and non-swappable address spaces.

System UIC	0-2	3-5	6-7	8	9	10-13	14	15+
Swappable (sec)	1	1	1	1	2	2	2	2
Non-swap (sec)	1	2	3	3	3	4	5	6

Figure 2-30. UIC Update Rates

For example, if the system UIC is six, then the UIC for non-swappable address spaces is updated every three seconds.

The UIC for swappable address spaces needs to be maintained more closely in order to effectively trim to the working set on a swap out. For non-swappable address spaces, it is not necessary to keep close track of the working set if the demand for real storage is light. This also gives some page stealing protection for non-swappable address spaces.

2.2.4.3.3 Page Steal Algorithm

The page replacement process in MVS maintains a queue of free pages to use in satisfying requests for new pages. This queue, called the Available Frame Queue (AFQ), is replenished by stealing pages from the address spaces in real memory. Page stealing is performed based on various values of the UIC of each page. Thus the UIC is used to define the working set of each address space. The changes described in this section were made in response to various operational problems that were discovered with continued field experience with MVS.

These changes were always made with two primary objectives in mind:

- o To achieve a balance holding enough page frames on the AFQ.
- o To distribute the burden of paging across the address spaces in real memory.

In MVS/XA and MVS/ESA systems, the page steal process takes four pages from each address space or common area before making a second pass to steal more pages.

The page stealing process is invoked when the available frame queue falls below a threshold and steals enough pages to maintain the queue at a target level. When the page stealing process is invoked, the RFR takes all pages whose UIC equals the system UIC. No differentiation is made between address spaces.

This algorithm corrects problems of the previous algorithms. It also tends to steal more pages from the types of address spaces that were previously favored, the low CPU users. Thus, the non-swappable online systems (e.g., CICS) that use their pages less frequently than other address spaces experience heavier paging rates.

2.2.4.4 Expanded Storage Concepts

This section introduces concepts common to expanded storage. Originally, this was a high-speed semiconductor memory, initially made available for the IBM 3090 processor series. In later systems, this became a redefinition of central storage performed during IPL. It was provided to maintain compatibility with existing applications. Additionally, with WLM Goal mode, many of the metrics discussed in the following sections and the control mechanisms used were eliminated. The information is presented here in a historical context only, and is no longer used for managing system performance.

The following sections discuss expanded storage and its impact on the job of Real Storage Management.

- 1 - Overview of Expanded Storage
- 2 - System Control of Expanded Storage
- 3 - Users of Expanded Storage
- 4 - Measurement, Analysis, and Management

2.2.4.4.1 Overview of Expanded Storage

Expanded storage is optional, high-speed semiconductor storage initially made available for the IBM 3090 processor series. The storage is installed in the central processor complex; that is, it is an integrated part of the complex and not a stand-alone piece of equipment.

In some respects, expanded storage can be thought of as a slower speed extension of central storage. The two differ in that expanded storage is not directly addressable and can be accessed only by moving 4K blocks of data from expanded storage into central storage. All data movement to and from expanded storage must be through central storage. There is no way to directly transfer data between expanded storage and auxiliary storage.

In other respects, expanded storage is similar to high performance paging devices such as the 2305-2 or 4305. However, the system control program knows more about expanded storage than it does about solid state paging devices. It uses this intelligence to optimize storage access, including paging and swapping operations and special kinds of application and system control data in ESA systems.

A primary use of expanded storage is as a substitute for auxiliary storage (page/swap data sets on DASD) in situations demanding better performance. The time it takes to read a page from auxiliary to central storage in a well-tuned system may range from a few milliseconds (with cached DASD controllers or solid-state channel-attached devices on dedicated paging channels) to 20 to 30 milliseconds for non-cached 3380 or 3390 DASD on non-dedicated channels. Indirect measurements suggest that typical times for page transfers from expanded to central storage take something less than 100 microseconds. This represents a reduction in page resolution time of approximately two orders of magnitude.

When expanded storage was first introduced, the system control program used it exclusively for paging and swapping. New MVS versions have introduced additional uses of expanded storage, such as virtual I/O (VIO) and hiperspaces.

Expanded storage, therefore, added a new level to the storage hierarchy between central and auxiliary storage. Now, page movement within the storage hierarchy can be divided into the

following categories:

From	To	Terminology
-----	-----	-----
Central	Auxiliary	Page-out or swap-out
Auxiliary	Central	Page-in or swap-in
Central	Expanded	Expanded storage write
Expanded	Central	Expanded storage read
Expanded	Auxiliary	Page migration

2.2.4.4.2 System Control of Expanded Storage

The Real Storage Manager (RSM) handles the transfer of pages between central and expanded storage. RSM determines both whether a page will be written to expanded storage, and when a page in expanded storage should be migrated to auxiliary storage. It bases these decisions on four values:

1. Criteria age

The installation-specified time value that is used by RSM to determine whether a page should be transferred from central to expanded storage. Criteria ages are specified for various categories of pages and workloads in the SYS1.PARMLIB member IEAOPTxx. IBM supplies default values that are reasonable for most normal processing environments.

2. System-high unreferenced interval count (UIC)

The time since the least recently referenced page in central storage was referenced. This value is therefore an inverse measure of central storage contention (a low UIC indicates a high demand for central storage). It is not possible to explicitly define a minimum value for the UIC, but it is influenced by several parameters in the IEAOPTxx member of SYS1.PARMLIB.

3. Migration age

The time that a page has remained unreferenced in expanded storage. It is an inverse measure of expanded storage contention (a low migration age indicates a high demand for expanded storage). In periods of contention, RSM moves pages in expanded storage that have remained unreferenced for this period of time into central storage, and the auxiliary storage manager (ASM) then pages them out to auxiliary storage. This process is referred to as page migration.

4. Think time

The average delay between the time when the system is ready for new input from a terminal user and the time when the terminal user actually provides such input by pressing the "enter" key or equivalent. This value is used primarily to control logical swapping, but is also used in terminal input and output wait swapping to expanded storage.

Criteria age, is specified for individual types of pages and address space by the installation (or allowed to default to IBM-specified values). The others are measured by the system control program and will vary with the nature and volume of work being performed and with system configuration. All are specified or measured in seconds.

In paging situations, RSM sends changed and unchanged stolen pages, virtual fetch pages, and page-out requested pages to expanded storage when the migration age is greater than the criteria age.

In swapping situations, RSM normally sends both swap-out trim and working set pages to expanded storage when the migration age is greater than the criteria age for the type of address space in question. However, TSO terminal wait, steal, and pageout users' working sets are sent to expanded storage when the sum of migration age and UIC is greater than the sum of criteria age plus think time.

When a page is required in central storage, RSM first checks to see if the page is in expanded storage. If not, a page fault is generated and the page is retrieved asynchronously from auxiliary storage by the ASM with a normal page-in operation. If the page is in expanded storage, however, no page fault is generated and RSM immediately transfers the page into central storage. This transfer happens synchronously with program execution. It is important to emphasize that a page fault is not generated in this case, and the expanded-to-central transfer is not considered a page-in. Indeed, RMF Version 3 does not even report page movement from expanded to central storage.

Swapping to expanded storage differs from swapping to auxiliary storage. When an address space is swapped out to expanded storage, RSM and the System Resources Manager (SRM) divide the working set pages into primary and secondary working set groups. The primary group, which is managed by RSM as a single entity, contains all LSQA and fixed pages, plus one page from each virtual storage segment included in the whole working set. The secondary group comprises all working set pages not included in the primary group.

Once the address space has been swapped out to expanded storage, RSM recognizes the difference between the primary and secondary working sets when it makes its decisions about migration. Secondary pages are migrated first, and only after all secondary pages have been migrated will the primary

working set be migrated.

If the address space is swapped back in before it has been migrated, the primary working set is swapped in first, at which time the address space becomes dispatchable. Secondary working set pages are then treated the same as all other pages in expanded storage. This means that if they are migrated, they will be sent to local page data sets rather than swap data sets. If you do a substantial amount of swapping to expanded storage, this could be an argument in favor of an auxiliary storage subsystem configuration consisting only of local page data sets with no swap data sets.

2.2.4.4.3 Users of Expanded Storage

The RSM categorizes expanded storage use by the type of paging group and by the status of the page being considered for placement in expanded storage. The paging group can be a:

Type 0 - Privileged or non-swappable (including all common areas) address space,

Type 1 - Any other type of page not included in types 0 or 2, or

Type 2 - A page belonging to a TS0 user waiting for terminal I/O to complete, or pages paged-out or stolen from a TS0 user.

The page's status is categorized as a:

Changed paged-out page
Changed swapped-out page
Changed stolen page
Unchanged paged-out page
Unchanged swapped-out page
Unchanged stolen page
Working set page ready for swap-out
Virtual fetch page
VIO swapped-out page
Hiperspace page

NOTE: Hiperspaces were introduced with MVS/ESA, and are not available in MVS/XA. The use of expanded storage for VIO was standardized in MVS/ESA, but is available in MVS/XA systems by means of applying IBM APAR YA09186.

Different criteria ages can be specified for each combination of paging type and page status, with the exceptions of virtual fetch, VIO, and hiperspace pages. The exceptions have a single criteria age specification without paging type classification differences. This allows the installation to tailor the configuration so that expanded storage usage is prioritized by the type of address spaces that are most important in the specific processor complex. It is possible to set the parameters so that some of the combinations are eligible and others ineligible to use expanded storage.

2.2.4.4.4 Measurement, Analysis and Management

RMF measures and reports on many data elements concerning paging, swapping, and memory use. Beginning with RMF 3.3.0 and continuing through RMF Version 4 (introduced with ESA), much information has been added that differentiates between paging or swapping to expanded versus auxiliary storage, and between the use of central versus expanded storage for CSA, LPA, SQA, etc., pages. Expanded storage has added to the complexity of understanding and managing storage hierarchy.

Even though well-tuned systems or systems with low storage contention are likely to experience expanded storage-related delays so small as to be relatively unimportant when compared to other sources of delay, installations with expanded storage should now monitor its utilization and activity as well as auxiliary storage activity. Expanded storage delay can become substantial to the point of significantly reducing the benefit from your investment in it. The System Storage Usage Report described in section 2.2.3.2 can aid in tracking the utilization of expanded storage.

In a system with expanded storage, the degradation due to page movement delay from expanded storage to central storage is not directly measurable. Indeed, the rate of movement from expanded to central is not measured by RMF in MVS/XA, it is only measured in MVS/ESA systems. While the system control program maintains (and RMF records) the average time that it takes to read a page from auxiliary to central storage, no such estimate is made of the time it takes to transfer a page from expanded to central storage.

Page migration should be a primary concern because, if excessive, it represents ineffective page placement and ineffective use of expanded storage. If there is contention for expanded storage, it makes no sense to move a page from central to expanded storage if that page will eventually be moved back into central and then paged out to auxiliary. It would be much better to page it out directly to auxiliary storage.

Minimizing such ineffective page placement depends on how accurately an address space's future activity and storage references can be predicted. RSM makes these predictions based on the criteria age values specified in the IEA0PTxx parameters for the various categories of expanded storage use.

Expanded storage can be used to improve the response time of

online systems, to increase the capacity of the processor complex, or a combination of the two. The best way to use expanded storage depends on the business and technical environment in which it is used.

Unfortunately, the mechanisms used to control expanded storage allocation do not provide for workload assignments, such as performance group or domain. It is therefore necessary to understand the storage utilization characteristics of the individual workloads and the system as a whole before attempting to modify the criteria ages that affect expanded storage control.

There are some general guidelines that apply to all systems, however. To improve the response times of online systems running as non-swappable address spaces, set low criteria ages for the type 0 paging group and high ages for types 1 and 2 paging groups. This favors the onlines over other work, but only if there is other work in the same system. In a system used for application development and testing with ISPF/PDF users submitting batch compiles and executions, you might want to favor type 2 over types 0 and 1. If you have a processor complex that is performing a wide variety of tasks and with changing workload characteristics over time, then to achieve both somewhat better response time and more throughput use the default values for all criteria ages.

When making criteria age changes, keep in mind the following:

- o Assigning a low criteria age to a particular page type and status is saying, in effect, that you want RSM to move such pages to expanded storage before other page type/status classifications even if the contention for expanded storage is high. You would do this if
 - (a) you regard the workload owning the page as important and deserving of a performance improvement, or
 - (b) you expect the page to be referenced quickly and want to lower the load on ASM by making it unnecessary to page-out and then page-in the same page in a brief time period.
- o Assigning a high criteria age to pages has the opposite effect: such pages are sent to expanded storage by RSM after page types with lower criteria ages. In periods of contention, such pages tend to go to auxiliary storage instead. You would do this if you either
 - (a) want to use your expanded storage resources for more

- important work, or
- (b) expect that such a page will not be referenced for a relatively long time and would be migrated to auxiliary storage anyway.

Although IBM has established reasonable default values, it is ultimately the responsibility of individual installations to set these values for best use of expanded storage. The performance analyst must monitor page movement to ensure that the specified values are optimizing the use of expanded storage for the installation's requirements. If criteria ages are set too low, expanded storage will tend to be over-utilized. This will result in a low migration age, large values for both central-to-expanded page movement, and page migration. If the difference between these two measurements is small, it could indicate that too many pages are being migrated.

2.2.5 Analytic Tutorial

This section is not available at this time.

2.2.6 Case Study

This section is not available at this time.

2.2.7 Component Operation

This Component Operation Section describes the operation of the Real Storage Analysis routines in terms of the CA MICS Performance Manager inquiry options selected.

Inquiry options available with the Real Storage Analysis software describe the reports to be generated. In particular, these options define:

- o The cycles from the CA MICS RMF or SMF files to be processed.
- o Workload selection by system, workload group, date/time, and zone.
- o The summarization level of the reports as hourly, daily, weekly, or monthly. The availability of each summarization level assumes that your system supports the required files at each level.
- o The reports and plots to be produced.

The following fields are common to most control panels:

Timespan:	The timespan (DETAIL, DAYS, WEEKS, or MONTHS) in the CA MICS files you want to use as input for processing.
Cycles:	The one or more cycles from the CA MICS database that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a CYCLE from the CA MICS database will usually contain a 24-hour period spanning two days. Therefore, some other type of selection may be required to generate a daily report. This selection capability is provided by the DATE and HOUR control options.
Database:	The single-character CA MICS identifier for the database to be accessed. Valid values are the database unit identifiers available at your installation. Contact your system administrator if you do not know what database units are available, or enter a question mark (?) to view a list of your database units. You can select up to five databases per inquiry.
Allow	The indication of whether the database

Runtime Override:	identifier and/or the cycles specification can be overridden at execution time.
Report by hour:	The indication of whether a report or plot should be generated for each hour.
Generate a report for each value of Zone:	The indication of whether a separate report or plot should be generated for each zone defined in the input data or whether all zones should be merged together for the report or plot.
Date:	Up to two date ranges (DDMONYY format) to be used for analysis. This option is normally used to drop the extra day from a cycle caused by the timing of the SMF dump.
Hour:	One to three ranges of hours for which data is to be reported. For an individual hour, leave the end hour blank or make it the same as the start hour. You can specify a value from 00 to 23.
Zone:	The specification that zone selection is to be applied. If zone 1 definition is 8:00 a.m. to 6:00 p.m., you can specify reports covering that time period by entering a 1 in the ZONE field.
SYSID:	The specification that SYSIDs are to be selected for reporting. You can enter up to six SYSIDs to be selected. If you leave the field blank, no testing of SYSIDs will be done.
SVC Class:	The WLM Service Classes to be selected for reporting. If you code this field, only those specified will be selected.
SVC Cls. Pd:	The WLM Service Class Periods to be selected for reporting. If you code this field, only those specified will be selected.
Job Name:	The specific JOBNAMEs to be reported. If you code this field, only those specified will be selected.
Workload Type:	The type of workload to be reported (such as Performance Group, Domain).
Workload Input	The type of file that will be used for reporting (such as Batch Program, TSO file, or Started

File: Task file).

Specify extended options: The specification that allows you to supply additional report specifications. These options include special SAS data selection statements, non-standard input and output files, and special graphics options.

The operation of the time selection control parameters depends on definitions made during the installation of the CA MICS system. If you are not sure about the ZONE definitions for your site or the CA MICS database update procedures, ask your system administrator.

Figure 2-58 shows the Real Storage Management (RSM) Analysis menu that you use to select the desired report or plot within the RSM reporting area.

```

/----- Real Storage Management (RSM) Analysis -----\
| Option ==>                                           |
| Composing CA MICS Inquiry: RSMT1 - Real Storage Manager Inquiry |
|                                                       |
|     1 - RSM System Storage Usage Report              |
|         Real storage usage report of system areas.   |
|     2 - RSM Private Storage Usage Report             |
|         Real storage usage report of private address spaces. |
|     3 - RSM Swap Analysis Report                    |
|         Analysis of Swap Reason or Swap Type.       |
|                                                       |
\-----/

```

Figure 2-58. RSM Analysis Selection Panel

The following sections describe the inquiry panels for each RSM report or plot:

- 1 - RSM System Storage Usage Report
- 2 - RSM Private Storage Usage Report
- 3 - RSM Swap Analysis Report

2.2.7.1 RSM System Storage Usage Report

The RSM System Storage Usage report details the use of real storage by the entire system. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the RSM System Storage Usage Report obtains its data from the SCPPAG Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The RSM System Storage Usage Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options that are normally defaulted.

```

/----- RSM Real Storage System Usage Report -----
| Command ==>
|
| Composing CA MICS Inquiry: STRK1 - System Storage Tracking
|
| Time-span ==> DETAIL (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 02
| Data base ==> A (PRIMARY)      _      _      _      _
| Allow Run-Time Override of Data Base ==> Y (Y/N) and/or Cycle(s) ==> Y (Y/N)
|
| Report by hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone      ==> N (Y/N)
|
| Data Selection:                                           Allow run-time override
|   Date      ==> _____ - _____ - _____ ==> N (Y/N)
|   Hour      ==> __ - __      __ - __      __ - __      ==> N (Y/N)
|   Zone      ==> _____ ==> N (Y/N)
|   Sysid     ==> _____ ==> N (Y/N)
|
|
| Specify extended options ==> N (Y/N)
|
/-----

```

Figure 2-59. RSM System Storage Area Usage Tracking Report Inquiry Panel

Figure 2-59 shows an RSM Real Storage System Usage Report inquiry panel. The fields common to all Real Storage Management Analysis routines are described in Section 2.2.7. The following are general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Report by hour and Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized to the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

There are four other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.2.7.2 RSM Private Storage Usage Report

The RSM Private Storage Usage Report details the use of real storage by WLM Service Classes or Service Class periods. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the RSM Private Storage Usage Report obtains its data from the SCPPAG, WLMSEC, or WLM_EC file, depending on the workload type you specify on the panel. If you are unsure of the timespans supported at your site, ask your system administrator.

The RSM Private Storage Usage Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options that are normally defaulted.

```

/----- RSM Private Storage Usage Report -----
|
| Command ==>
|
| Composing CA MICS Inquiry: PRIV1 - Private Area Tracking
|
| Timespan ==> DETAIL (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 02
| Database ==> A (PRIMARY) _ _ _ _ _
| Allow Run-Time Override of Database ==> Y (Y/N) and/or Cycle(s) == > Y (Y/N)
|
| Report by hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone ==> N (Y/N)
|
| Workload Input file ==> PGA (DMA/PPA/PGA/RG) Numeric Element ==> PGAAVAFT
|
| Data Selection:                                         Allow run-time override
| Date ==> _ _ _ _ _ - _ _ _ _ _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N)
| Hour ==> _ _ - _ _ _ _ - _ _ _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N)
| Zone ==> _ _ - _ _ _ _ - _ _ _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N)
| Sysid ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> N (Y/N)
| SVC. Class ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _
| SVC. Cls. Pd ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _
|
| Specify extended options ==> N (Y/N)
|
\-----

```

Figure 2-60. RSM Private Storage Usage Report Inquiry Panel

Figure 2-60 shows an RSM Private Storage System Usage Report inquiry panel. The fields common to all RSM Analysis routines are described in Section 2.2.7. The following are

general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization to the ZONE level can be done for any timespan. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Report by Hour and the Generate a Report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized into the timespan level (that is, MONTHS is summarized to the month, WEEKS to the week, and so on).

The numeric element for which reporting is done defaults to fffAVAFT, where fff varies according to the input file being used. You can replace this value with the name of any numeric element present on the file you are using. Enter a question mark (?) in this field to obtain a list of all the elements on the selected file. Prior to RMF 3.5, fffAVRSF can be used for analyzing storage usage instead of fffAVAFT, which first became available with RMF 3.5.

Not all data selection values are valid for every type of input file, because some of the elements being tested do not exist in certain files. An error message is displayed on the panel if you specify an invalid selection.

There are four other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus, see Section 1.4.3.6, Specifying Common Extended Options.

2.2.7.3 RSM Swap Analysis Report

The RSM Swap Analysis Report lists measurements from the CA MICS SCPSWP File(s) presented by swap reason or by swap type. Data can be selected by system, date, and/or time period.

The Swap Analysis Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional parameters that are normally defaulted.

```

/----- RSM Swap Analysis Report -----\
| Command ==>
|
| Composing CA MICS Inquiry: KLSX - SWAP REPORT
|
| Time-span ==> DAYS (DETAIL/DAYS/WEEKS/MONTHS) Cycle(s) ==> 01 - 02
| Data base ==> A (PRIMARY)
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N)
|
| Report by hour if timespan is DAYS or WEEKS ==> Y (Y/N)
| Generate a report for each value of Zone ==> Y (Y/N)
|
| Report swap reason or swap type ==> REASON (REASON/TYPE)
|
| Data Selection: Allow run-time override
| Date ==> - - - - - ==> N (Y/N)
| Hour ==> - - - - - ==> N (Y/N)
| Zone ==> - - - - - ==> N (Y/N)
| Sysid ==> - - - - - ==> N (Y/N)
|
| Specify extended options ==> N (Y/N)
|
\-----/

```

Figure 2-61. RSM Swap Analysis Report Inquiry Panel

Figure 2-61 shows an RSM Swap Analysis Report inquiry panel. The fields common to all RSM Analysis routines are described in Section 2.2.7. The following are general comments on this panel.

You can specify the organization of the report as either REASON or TYPE (REASON is the default). When you specify

REASON, each report category represents a reason and the total swaps by type are line items. When you specify report by TYPE, each category represents the swap type and line items identify total swaps by reason.

There are four other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.2.8 Requirements

The Real Storage Management Analysis reports let you track the allocation of real storage on a monthly, weekly, daily, or hourly level. Paging and swapping workloads are reported to allow analysis of the relation between real storage allocation and the operation of the page/swap subsystem. Real storage requirements at the performance group level are reported to help you in evaluating the impact of new applications or users, as well as the growth of existing applications. Swapping statistics are reported to assist you in investigating swapping overhead.

Figure 2-66 shows the CA MICS files and timespan levels that must be maintained at your installation in order to produce all of the Real Storage Management Analysis reports.

CA MICS FILE	LEVEL				
	X	D	W	M	Y
SCPPAG	x	x	x	x	
SCPSWP	x	x	x	x	
WLMSEC	x	x	x	x	
WLM_EC	x	x	x	x	
BATPGM	x				
BAT_ST	x				
BAT_TS	x				

Figure 2-66. RSM Analysis CA MICS File Requirements

The following is a list of files required to run each individual report:

- o System Storage Usage Report - SCPPAG
- o Private Storage Usage Report - WLMSEC or WLM_EC
- o Swap Analysis Report - SCPSWP

Only one level is required for the CA MICS RMF product files (for example, DETAIL, DAYS, WEEKS, or MONTHS), depending on the timespan you specify during inquiry definition.

The number of cycles that are maintained at your site for

each of the CA MICS files in Figure 2-66 determines the time frame for which you can produce RSM Analysis reports and plots.

2.3 Auxiliary Storage Management Analysis

Virtual systems must maintain a copy of all virtual memory pages. In MVS, this function is carried out by the Auxiliary Storage Manager (ASM). This chapter describes the structure and operation of ASM. Auxiliary storage is defined as several direct access data sets, each of which provides storage capacity for one of two major areas:

- o Paging space
- o Swapping space

In MVS, ASM supports three logically distinct groups of pages: pages from the pageable link pack area, pages from the common system area, and pages from the private area. At least one page data set must be defined for each group of pages.

There is a single page data set to hold pages from the pageable link pack area (the PLPA data set), one to hold common system area pages (the common data set), and one or more to hold pages from the private area (the local data sets). Organizing auxiliary storage in this manner allows I/O overhead to be spread across several data sets and permits concurrent processing of many virtual pages.

This chapter describes the use of these page and swap data sets by ASM, including the paging control algorithms, data set selection algorithms, and the swapping process. ASM operation is described through MVS/SP 2.1.2 (MVS/XA). Data set configurations and device selection considerations in using drums and 3330, 3350, and 3380 devices are discussed.

The reports produced by the CA MICS Auxiliary Storage Management Analysis inquiries give you the capability to monitor the operation of ASM in terms of utilization and response by:

- o Individual page or swap data sets.
- o Page space type (e.g., the local data sets).
- o Device type (to compare the service provided by 3380s to that provided by 3350s, for example).

You can produce these reports from the CA MICS RMF product data at any supported timespan and summarization level using the SCPPSDnn and HARDVAnn Files. The availability of each summary option depends on the generation of the SCPPSD Files

at the corresponding level of summarization. The CA MICS default is to produce files at the DETAIL, DAYS, and WEEKS levels, but not at the MONTHS level.

The response time reports indicate problems that result from the impact of paging I/O on the I/O subsystem or the impact of the I/O configuration on the page/swap data sets. You can use the reports described in Section 2.4, I/O Configuration Analysis, to examine these problems.

The following sections describe the operation of the Auxiliary Storage Management Analysis inquiries and how to use the reports that are produced:

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output
- 4 - MVS Concepts
- 5 - Case Study
- 6 - Component Operation
- 7 - Requirements

2.3.1 Functional Description

The Auxiliary Storage Management (ASM) Analysis inquiries allow you to track the utilization of the page and swap data sets to find problem areas in the ASM configuration.

Three basic report formats provide a view of the utilization and performance of the data sets comprising the ASM auxiliary storage configuration. The reports summarize this data by page space type (for example, for all local data sets), by device type, and by individual data set name.

This gives you the capability to identify parts of the ASM configuration that might be causing response problems due to the placement of a single data set, to measure the impact of not isolating an ASM data set, or to quantify the expected benefits from replacing 3350 devices with 3380 devices or a drum (that is, a fixed head device).

The inquiry routines use data from the page data set files in the RMF SCP Information Area, combined with the device data from the RMF HAR Information Area, to estimate the elongation of the page or swap service times due to device delays. Figure 2-57 illustrates data that summarizes the utilization and response times of the ASM data sets on 3380 devices.

CA MICS Performance Manager DETAIL ASM DEVICE-TYPE ACTIVITY REPORT SYSTEM = TS04 RUNDATE = JAN 19, yyyy DEVICE TYPE = 3380											

INTVL/DAY (HR:MM-DD):	8:13-13	8:29-13	8:43-13	8:59-13	9:13-13	9:29-13	9:44-13	9:58-13	10:14-13	10:29-13	CHANGE %

DATASET SIZE											

AVG SLOTS ALLOCATED	25688	25688	25688	25688	25688	25688	25688	25688	25688	25688	---
AVG SLOTS USED	3559.25	4103.00	4460.00	4778.00	5135.00	5639.00	6113.00	6173.00	6373.00	6526.00	83.35 %
MAX SLOTS USED	6758.00	7675.00	8435.00	8937.00	10017	11660	11892	11960	12328	12856	90.23 %
UNUSABLE SLOTS	0	0	0	0	0	0	0	0	0	0	N/A
PERCENT SLOTS USED	13.85 %	15.97 %	17.36 %	18.60 %	19.99 %	21.95 %	23.79 %	24.03 %	24.80 %	25.40 %	83.33 %

DATASET WORKLOAD											

TOTAL PAGING I/O	2741	7289	5394	6846	5829	12817	13017	16346	19069	16540	503.42 %
TOTAL DEVICE I/O	2842	8204	5689	7527	6386	14659	15200	19305	23535	19880	599.50 %
I/O CONTENTION INDEX	0.037	0.126	0.055	0.099	0.096	0.144	0.168	0.181	0.234	0.202	445.94 %
AVG SIO RATE	0.76	2.02	1.50	1.90	1.62	3.56	3.62	4.54	5.30	4.59	503.94 %
MAX SIO RATE	1.53	3.62	2.93	3.51	2.96	6.43	6.22	7.48	8.33	7.51	390.84 %
AVG PAGE RATE	8.85	16.90	17.05	15.96	15.16	27.92	29.75	36.81	50.03	47.01	431.18 %
AVG PAGE BURST SIZE	11.62	8.35	11.38	8.39	9.36	7.84	8.23	8.11	9.45	10.25	-11.79 %
AVG DATASET UTIL	3.19 %	7.91 %	5.97 %	6.22 %	5.69 %	13.15 %	13.66 %	16.87 %	25.58 %	21.94 %	588.08 %
MAX DATASET UTIL	6.44 %	15.22 %	11.66 %	12.55 %	11.22 %	26.85 %	26.00 %	31.81 %	47.27 %	42.06 %	553.10 %

DATASET SERVICE											

AVG RESPONSE TIME/IO	0.042	0.039	0.040	0.033	0.035	0.037	0.038	0.037	0.049	0.048	14.28 %
AVG SERVICE TIME/IO	0.042	0.039	0.040	0.033	0.035	0.037	0.038	0.037	0.048	0.048	14.28 %
AVG SERVICE TIME/PAGE	0.004	0.005	0.004	0.004	0.004	0.005	0.005	0.005	0.005	0.005	25.00 %
MAX SERVICE TIME	0.004	0.022	0.043	0.013	0.016	0.017	0.022	0.020	0.020	0.017	325.00 %

Figure 2-67. Sample ASM Activity Report by Device Type

2.3.2 Usage Guidelines

The Auxiliary Storage Management Analysis inquiries provide two basic reports that use data from the MVS Page/Swap Data Set Activity File (SCPPSD) and the Device Activity File (HARDVA) to report on the utilization and overall response given by the ASM configuration. These reports allow tracking of the ASM configuration on an hourly, daily, weekly, or monthly basis if the SCPPSD and HARDVA Files are supported in the CA MICS database at the appropriate levels of summarization.

This section applies the ideas discussed in Section 2.3.4, MVS Concepts, to these environments, discusses minimal ASM configurations, and gives some basic recommendations. As with all recommendations, these must be considered as they apply to your specific environment. If you are not familiar with the terminology and concepts associated with the Auxiliary Storage Manager, you may find it helpful to read Section 2.3.4 before continuing with this section.

All the recommendations given here or in any of the references are to be considered only as they apply to your environment, not as strict rules.

To aid in evaluating the response and utilization levels of the ASM configuration, the Auxiliary Storage Management Analysis reports allow data describing the ASM configuration to be viewed for individual page or swap data sets, by data set class (e.g., pageable link pack area (PLPA), common, local, or duplex), or by device type (e.g., 3380s).

The following sections present ASM configuration guidelines, an example on sizing data sets, and considerations on using specific device types.

- 1 - ASM Configuration Tuning
- 2 - ASM Data Set Sizing

2.3.2.1 ASM Configuration Tuning

The interaction of the ASM configuration with the I/O subsystem is not directly measured by the Auxiliary Storage Management Analysis reports. The ASM Page/Swap Response Time Report shows if service times are uniform or if the response of specific data sets is causing a problem. The ASM Page/Swap Activity Report can be used to evaluate the impact of not isolating page or swap data sets on dedicated volumes. Both data set response and service times are reported.

The following list gives some configuration guidelines that are generally applicable to all levels of ASM support:

1. Three data sets is the minimum paging configuration. The first data set defined is for PLPA, the second for common, and the remaining data set(s) are for local data sets. A minimum paging configuration of only three page data sets (i.e., one PLPA, one common, one local) is not very practical. Check page rate by device and data set response times to see if more local page data sets are required.
2. In pre-MVS/XA systems, use swap data sets in a TS0 environment or for an environment with heavy swapping. Overflow of swapping into the local data sets adversely affects the performance of the local data sets. In MVS/XA or MVS/ESA systems, because of improvements in the paging data set algorithms (as opposed to swap data set algorithms), you should strongly consider the use of local page data sets for swapping.
3. Define extra page and swap data sets (if used). Use the PAGENUM keyword in the IEASYSnn member of SYS1.PARMLIB to allow for the addition of extra data sets that can be used to recover from an emergency condition that exhausts auxiliary storage.
4. If the PLPA paging rate is very high, PLPA can be allowed to spill into the common data set. The initial configuration should be sized in accordance with information presented in the IBM Initialization and Tuning Guides. The WEEKS level of the ASM Page/Swap Activity Report provides the information necessary to make this decision.

If you use intentional PLPA overflow to common, be careful to make the common data set large enough to hold both the PLPA overflow and all common system area pages.

If the common data set is filled, the ASM stops using it and the system enters a wait state.

5. Do not put ASM data sets on channels with non-RPS (Rotational Position Sensing) activity. The non-RPS activity will tend to dominate, because the ASM releases the channel during sector searches. This condition may be recognized by a build-up of the ASM queue.
6. If fixed head devices are available, allocate them to the highest activity data sets. Most load modules are normally placed in the link pack area because they are frequently referenced, which means the PLPA data set is a good candidate for a fixed head device.
7. Avoid placing more than one ASM data set on any single moveable head device. ASM should be prevented from competing with itself at the device level.
8. Avoid placing ASM data sets on any shared device. The slot selection algorithms all depend on knowledge of the arm position following the last ASM I/O request. ASM channel programs cannot operate efficiently when another CPU has moved the arm.
9. Over specify paging space. Larger common and local data sets on moveable head devices increase the probability that the slot selection algorithms will find slots underneath or near the current arm position. Thus, backing store requirements will not inhibit address space creation or the allocation of virtual I/O (VIO) data sets.
10. Use multiple local and swap data sets (if swap data sets are used at all). Even when paging space or swapping space requirements can be met by a single data set, multiple data sets allow ASM to service multiple requests concurrently. This is especially important during peak loads.
11. Duplex Data Sets. The use of a duplex data set severely affects common area page-out processing. You should perform a periodic evaluation of the actual saving (i.e., how many IPLs were saved) to determine if the savings is worth the overhead.
12. Allocate entire volumes for ASM data sets, and place no more than one ASM data set on a single path.

13. If real storage is constrained, reduce the amount of logical swapping by lowering the maximum system think time value on the LSCTMTE parameter in member IEAOPTX of SYS1.PARMLIB. See the Initialization and Tuning Guide for a description of this parameter.

2.3.2.2 ASM Data Set Sizing

The ASM Page/Swap Activity Report shows the slots allocated and the average and maximum number of slots that were used. If the SCPPSD and HARDVA Files are available at the WEEKS or MONTHS level, you can use this report directly to adjust the paging and swapping space data set sizes. While excessive allocation is wasteful of real storage, the data sets must be large enough to gain the benefit of the new slot allocation algorithms.

The Initialization and Tuning Guides give examples of selecting initial sizes for the ASM data sets. Section 2.3.4.2 presents techniques for sizing paging space data sets and Section 2.3.4.3 presents techniques for sizing swapping space data sets.

2.3.3 Standard Output

The Auxiliary Storage Management reports give you the capability to track the allocation of auxiliary storage at a monthly, weekly, daily, or hourly level. The reports use data from the page/swap data set files and the device activity files from the RMF SCP and HAR Information Areas to report on the activity levels and response of the ASM subsystem. The device activity files are used to obtain the I/O queue time for each of the page/swap devices, because the I/O service time in the page/swap files includes only the I/O transfer time.

To provide different views of the ASM information, the reports summarize the ASM information by page space type (common, pageable link pack area (PLPA), local, duplex, and swap) device type, or data set name.

The following sections describe the report formats of the ASM Page/Swap Activity and the ASM Page/Swap Response Time Reports:

- 1 - Common Report Elements
- 2 - ASM Page/Swap Activity Report
- 3 - ASM Page/Swap Response Time Report

2.3.3.1 Common Report Elements

This section describes the elements common to all the Auxiliary Storage Management reports (see Figures 2-68 to 2-71). Options that are unique to a report are documented with that report.

TITLE ELEMENTS

Each report has up to five lines of title information giving the installation name, component name, name of the report, run date, and the CA MICS SYSIDs being reported.

The headings on the ASM Page/Swap Activity and ASM Page/Swap Response Time Reports detail the data selected, the summary level of the report, and the time periods reported by column. For the monthly report, the column title is month/year; for the weekly report, the first day of the week is given; and for the daily report, the day is given as month/day/year. For the detail report, the hour:minute and day are given. For the hourly level, the hour and day of the month are given as the column title. The following elements are in the standard title part of each report:

Installation Name:	The installation name that is defined via panel at CA MICS system installation time. This value is centered in the first line of the report. For more information about this value, consult your CA MICS System Administrator.
Component Name:	The name of the performance component, centered in the second line of the report.
Report Name:	The name of the report, centered in the third line of the heading. It cannot be modified.
SYSID:	The CA MICS logical SYSID being reported. Data from multiple systems cannot be combined using the standard control parameters, although separate reports are provided for each system.
ZONE:	If data is being reported by zone, the zone number. If the zone option is not specified, the information for all selected zones is summarized, and no zone value appears on the report.
RUNDATE:	The report execution date in SAS WORDDATE format.

ADDITIONAL ELEMENTS

Time Interval:	The time intervals that are being reported listed on the first row of the report. This varies depending on the reporting options and the timespan being used.
%CHANGE:	The percent change for each report item (row) between the base period (first data column) and the last period (last data column), shown in the last column of the report. If data for only one period is reported, then dashes (---) are displayed. If the base period is 0, then N/A is printed to indicate a percent change could not be calculated.

2.3.3.2 ASM Page/Swap Activity Report

The ASM Page/Swap Activity Report tracks utilization and response of each page data set found in the RMF SCPPSD Files over the selected time range. A sample report is shown in Figure 2-68.

REPORT OPTIONS

The ASM Page/Swap Activity Report supports time, date, zone, and SYSID selection. The report options determine the level of the CA MICS database from which the data is selected, and the "by" value that is used for summarization.

DATA ELEMENTS

The elements reported are divided into three categories:

- o Data Set Size
- o Data Set Workload
- o Data Set Service

DATA SET SIZE

The data set size elements describe the allocation of the data set or data sets being described. When reporting by device type or page data set type, the elements in this section represent a total allocation over the data sets in a group (that is, over all local page data sets).

AVG SLOTS ALLOCATED:	The average number of slots (or swap sets) allocated to the data set during the reporting period. When reporting by device type or page data set type, this value presents an average profile for all the data sets in this group.
AVG SLOTS USED:	The average number of page data set slots (or swap data set swap sets) in use during this reporting interval. This value comes from the PSDAVU element in the SCPPSD File. RMF calculates this by sampling the Page Activity Reference or Swap Activity Reference Table each RMF cycle. When reporting by device type or page data type, this value presents an average profile for all the data sets in this group.
MAX SLOTS USED:	The maximum number of page slots (or swap sets) found to be in use on this data set during any

RMF sample. Note that the true maximum value can occur in between RMF cycles, but it cannot be smaller than this value. When reporting by device type or page data set type, this value is the maximum value found in any of the data sets included in this total. Investigations in sizing ASM data sets should always be done using this report summarized by data set name.

UNUSABLE SLOTS: The number of unusable slots (or swap sets) in this data set over the last RMF recording interval of the reporting period.

PERCENT SLOTS USED: The percentage of slots used over the reporting period. This value is equal to AVG SLOTS USED divided by AVG SLOTS ALLOCATED, multiplied by 100.

DATA SET WORKLOAD

The data set workload elements describe the work done by a page/swap data set or by the group of data sets represented by this report. You can use these elements to identify situations where more devices are needed, as well as situations where the device selection algorithm may be causing a problem. Values in this section represent an individual data set or a group of data sets, depending on the reporting option selected.

TOTAL PAGING I/O: The number of paging I/O requests directed to this page or swap data set or group of page or swap data sets. Obtained from field IORSION of the I/O request blocks (IORBs) for this page/swap data set. Recorded for MVS/SE and later systems only. For MVS/SP 1.3, this value includes Resume I/Os (RIOs).

TOTAL DEVICE I/O: The total number of all I/O requests (both paging and non-paging) directed towards this device or group of devices. Obtained from adding together the I/O counts in the HARDVA File for this device or group devices.

I/O CONTENTION INDEX: The ratio of the previous two values (paging I/O versus device I/O) for this device or group of devices. A high value could indicate a possible problem caused by having non-paging and paging data sets sharing the same device.

AVG I/O RATE: The average Start I/O rate per second for this data set or group of data sets during the RMF reporting interval. This value is only applicable to SE level MVS and above. For MVS/SP 1.3, this value includes Resume I/Os (RIOs). This value is calculated by CA MICS as:

$$\text{PSDPSIOR} = \text{PSDSIO} / \text{DURATION}$$

MAX I/O RATE: The highest rate of I/O request activity to this data set or group of data sets found for any RMF recording interval included in the reporting time interval. Recorded for MVS/SE and later systems only.

AVG PAGE RATE: The number of pages transferred per second. This element is applicable to "page" data sets only.

AVG PAGE BURST SIZE: The number of pages transferred per I/O. This element is applicable to "page" data sets only.

AVG DATASET UTIL: The percentage of time that the ASM (Auxiliary Storage Manager) considered this data set to be in use. RMF considers the data set to be in use if any IORB (I/O Request Block) associated with the data set is marked as being in use (data sets on multiple-exposure devices have multiple IORBs). For the device type or page data set type options, this is the average utilization of all data sets in the group. This value is calculated by CA MICS as:

$$\text{PSDPCUSE} = (100 * \text{PSDUSE}) / \text{SAMPLES}$$

MAX DATASET UTIL: The highest percentage "busy" encountered for this data set or group of data sets during any RMF recording interval included in the reported time interval. NOTE: This gives ASM's view of the data set. IOS's view may be different. Check the HARDVA Files for device busy. An indication of this difference will be given by the difference in the device service time and the device response time, which is defined below.

DATA SET SERVICE

The data set service elements give a measure of the average

time to handle I/O requests to the page/swap data sets. The device activity file is used to factor in the device delays to compute the response time element. The difference between these two elements gives an indication of shared volume problems.

NOTE: If you request reports by hour and are using the WEEKS timespan, the HARDVA File will be tested to make sure the file is in sequence by hour. If the file is not in sequence by hour, a warning message will be issued and several elements will be changed on the report to missing values.

AVG RESPONSE TIME/I/O:	The number of seconds to handle a page or swap request, including IOS queuing delays. This field is calculated by including the I/O queue delay from the appropriate device activity files.

AVG SERVICE TIME/I/O:	The number of seconds required to complete an average I/O to this page (or swap) data set. Maintained for MVS/SE or later versions only.

NOTE: Any contention that may cause IOS to queue ASM's request is not visible to ASM.

AVG SERVICE TIME/PAGE:	The number of seconds required to transfer a page. This element is applicable only to true "page" data sets (that is, not swap).

NOTE: Any contention that may cause IOS to queue ASM's request is not visible to ASM.

MAX SERVICE TIME:	The highest (worst) average service time for I/O requests found in any RMF recording interval in the time covered by this record.

CA MICS Performance Manager											
HOURLY ASM DATASET ACTIVITY REPORT											
FOR ZONE = 1											
SYSTEM = XA01 RUNDATE = SEP 19, yyyy											
DSN = SYS1.PAGE109 UNIT = 272 DEVICE TYPE = 3380 PAGE SPACE TYPE = LOCAL											

HOUR/DATE (HR-MM/DD): 00-09/yy 01-09/yy 02-09/yy 03-09/yy 04-09/yy 06-09/yy 07-09/yy 08-09/yy 09-09/yy 10-09/yy CHANGE %											

DATASET SIZE											

AVG SLOTS ALLOCATED	45000	45000	45000	45000	45000	45000	45000	45000	45000	45000	---
AVG SLOTS USED	4322.00	4213.00	3927.25	2935.00	2834.06	158.00	1186.00	2017.44	2166.50	2945.13	-31.85 %
MAX SLOTS USED	4657.00	4266.00	4197.00	3609.00	3009.00	334.00	1722.00	2240.00	2401.00	3373.00	-27.57 %
UNUSABLE SLOTS	0	0	0	0	0	0	0	0	0	0	N/A
PERCENT SLOTS USED	9.60 %	9.36 %	8.72 %	6.52 %	6.29 %	0.35 %	2.63 %	4.48 %	4.81 %	6.54 %	-31.87 %

DATASET WORKLOAD											

TOTAL PAGING I/O	21	358	331	276	252	29	443	1588	847	2445	1542.85 %
TOTAL DEVICE I/O	23	408	345	356	254	29	446	1602	849	2447	539.13 %
I/O CONTENTION INDEX	0.095	0.014	0.004	0.290	0.008	0.000	0.007	0.009	0.002	0.001	-98.94 %
AVG SIO RATE	0.37	0.10	0.09	0.08	0.08	0.05	0.12	0.44	0.24	0.54	45.94 %
MAX SIO RATE	0.66	0.15	0.15	0.26	0.11	0.05	0.16	0.77	0.28	0.68	3.03 %
AVG PAGE RATE	0.81	0.17	0.25	0.12	0.16	0.61	0.62	1.20	0.77	1.43	76.54 %
AVG PAGE BURST SIZE	2.21	1.68	2.70	1.53	2.12	12.79	5.02	2.72	3.29	2.64	19.45 %
AVG DATASET UTIL	0.80 %	0.13 %	0.27 %	0.02 %	0.06 %	0.32 %	0.22 %	0.55 %	0.58 %	1.00 %	23.45 %
MAX DATASET UTIL	1.44 %	0.33 %	0.55 %	0.11 %	0.22 %	0.32 %	0.44 %	1.10 %	0.66 %	1.44 %	---

DATASET SERVICE											

AVG RESPONSE TIME/I/O	0.022	0.014	0.030	0.004	0.008	0.069	0.018	0.013	0.025	0.018	-18.18 %
AVG SERVICE TIME/I/O	0.022	0.014	0.030	0.004	0.008	0.069	0.018	0.013	0.025	0.018	-18.18 %
AVG SERVICE TIME/PAGE	0.010	0.008	0.011	0.002	0.004	0.005	0.004	0.005	0.008	0.007	-30.00 %
MAX SERVICE TIME	0.011	0.014	0.015	0.003	0.016	0.005	0.005	0.006	0.009	0.009	-18.18 %

Figure 2-68. Sample Hourly ASM Page/Swap Activity Report

2.3.3.3 ASM Page/Swap Response Time Report

The ASM Page/Swap Response Time Report tracks the response of the ASM configuration by data set, page space type, or device type. You can produce the ASM Page/Swap Response Time Report at any summary level, with a maximum of ten reporting periods on a page. This report processes the data from the SCPPSD and the HARDVA Files.

NOTE: If you are reporting by hour and using the WEEKS timespan, the software validates that the HARDVA File is in sequence by hour. If it is not, the report is still produced, but it reports service time instead of response time.

REPORT OPTIONS

The ASM Page/Swap Response Time Report supports time, date, zone, and SYSID selection. The report options control the level of the CA MICS database from which the data is selected and the "by" value that is used for summarization.

DATA ELEMENTS

The ASM Page/Swap Response Time Report shows the average response time of each ASM data set (or group of data sets) found in the SCPPSD Files over the selected cycles of the CA MICS database. This report gives you a concise picture of the overall ASM configuration and its performance. To help you in evaluating the response figures, the device type of each data set is given if the name of the data set is small enough.

Figures 2-69 through 2-71 show the three forms of the ASM Page/Swap Response Time Report. The first example, in Figure 2-69, is the report produced by selecting the Page/Swap data set names report level option on the ASM Page/Swap Activity Report inquiry panel (see Figure 2-77 for a sample of the panel). This option produced reports on all data sets at the HOURS level.

Figure 2-70 shows the report produced by selecting the Page/Swap data set types report level option. Information at the data set level is summed by page space type (PSDPST in the SCPPSD File).

The final report, shown in Figure 2-71, shows the device type form of the report at the hourly level of summarization. For this report, the Page/Swap device types report level option

was selected.

HOURLY ASM DATASET RESPONSE REPORT											
FOR ZONE = 1											
SYSTEM = XA01 RUNDATE = SEP 19, yyyy											
HOURLY DATE (HR-MM/DD):	11-09/yy	12-09/yy	13-09/yy	14-09/yy	15-09/yy	16-09/yy	17-09/yy	18-09/yy	19-09/yy	20-09/yy	CHANGE %
COMMON DATASET											
SYS1.PAGE101C (3380)	0.000	0.000	0.000	0.000	0.000	0.000	0.022	0.024	0.000	0.012	N/A
LOCAL DATASET											
SYS1.PAGE109 (3380)	0.005	0.006	0.005	0.008	0.003	0.006	0.007	0.003	0.002	0.007	40.00 %
PLPA DATASET											
SYS1.PAGE101P (3380)	0.000	0.000	0.000	0.000	0.000	0.050	0.000	0.014	0.013	0.011	N/A
SWAP DATASET											
SYS1.SWAP102 (3380)	0.038	0.035	0.034	0.036	0.046	0.039	0.034	0.041	0.040	0.042	10.52 %

Figure 2-69. Sample ASM Page/Swap Response Time Report by Data Set Name

HOURLY ASM PAGE-SPACE RESPONSE REPORT SYSTEM = XA02 RUNDATE = SEP 19, yyyy											
HOUR/DATE (HR-MM/DD):	10-09/06	11-09/06	12-09/06	13-09/06	14-09/06	15-09/06	16-09/06	17-09/06	18-09/06	19-09/06	CHANGE %
COMMON DATASET	0.029	0.020	0.019	0.023	0.000	0.000	0.000	0.000	0.029	0.000	N/A
LOCAL DATASET	0.007	0.007	0.005	0.006	0.007	0.006	0.004	0.009	0.008	0.005	-28.57 %
PLPA DATASET	0.019	0.014	0.014	0.006	0.000	0.020	0.000	0.017	0.032	0.000	N/A
SWAP DATASET	0.036	0.018	0.031	0.050	0.031	0.033	0.044	0.022	0.032	0.034	-5.55 %

Figure 2-70. Sample ASM Page/Swap Response Time Report by Data Set Type

HOURLY ASM DEVICE-TYPE RESPONSE REPORT SYSTEM = XA01 RUNDATE = SEP 20, yyyy											
HOUR/DATE (HR-MM/DD):	00-09/06	01-09/06	02-09/06	03-09/06	04-09/06	05-09/06	06-09/06	07-09/06	08-09/06	09-09/06	CHANGE %
2305-2	0.000	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.007	0.005	N/A
3380	0.029	0.032	0.025	0.032	0.025	0.030	0.026	0.028	0.024	0.026	-10.34 %

Figure 2-71. Sample ASM Page/Swap Response Time Report by Device Type

2.3.4 MVS Concepts

The Auxiliary Storage Manager (ASM) maintains copies of all virtual storage pages that are not currently assigned to a real memory address. ASM controls two types of data sets:

- o Paging space data sets
- o Swapping space data sets

To provide as much parallel operation as possible and to spread out the I/O overhead associated with paging, the following classes of page data sets are supported by MVS:

- o PLPA -- Provides primary storage for paged-out portions of pageable link pack area modules.
- o Common -- Provides primary storage for paged-out portions of the common system area.
- o Duplex -- When duplexing is specified, a secondary copy of common system area and PLPA pages is maintained in the duplex data set. This provides for additional reliability. Initially, the PLPA is written to both the duplex and PLPA data set. Subsequently, all common system area page-outs are written to both the duplex and common data set. When an unrecoverable read error occurs in the common or PLPA data set, ASM attempts to read that page from the duplex data set.
- o Local -- The local page data sets hold all the paged-out private area virtual storage of individual address spaces that are in real memory (i.e., swapped-in), as well as most of the pages of swapped-out address spaces. Precisely which private area pages are directed to paging space data sets and which are directed to swapping space data sets during a swap varies with the different levels of the Auxiliary Storage Manager. Section 2.3.4.3.1., Physical Swap Algorithm, discusses the swap process for MVS/SP 1.3 and MVS/XA. Virtual I/O (VIO) data sets (special data sets accessed by applications that use paging to provide I/O functions) are also contained in the local data sets.

Swap data sets may be provided to speed up swap processing. They contain all local system queue area (LSQA) private area virtual pages and, when the extended swap IUP or MVS/SP 1.3 (or beyond) is used, the swap-in working set of all swapped-

out address spaces.

If swap data sets are not defined, LSQA pages (and any other pages directed to the swap data sets) are written to the local page data sets. This normally interferes with the demand paging I/O and, depending on the level of the system, may cause response degradation. In an MVS/XA system with a well-configured paging subsystem using 3380s, the amount of degradation is probably very small, if any, depending on the amount of swapping.

Section 2.3.2, Usage Guidelines, discusses tuning ASM configurations for the different paging and swapping algorithms and explains the use of the Auxiliary Storage Management reports in the analysis of the ASM configuration.

This concepts section discusses the swapping and paging algorithms used by ASM. The following sections explain ASM operation, the management of page traffic, and the principles of tuning an ASM configuration:

- 1 - Auxiliary Storage Management Overview
- 2 - Paging
- 3 - Swapping

2.3.4.1 Auxiliary Storage Management Overview

The purpose of the Auxiliary Storage Manager is to move virtual memory pages between real storage and auxiliary storage. The Auxiliary Storage Manager divides virtual memory pages into three logically distinct groups: PLPA, common, and local. At least one page data set is required for each group of pages in a minimal ASM configuration, though there are multiple local page data sets in most systems. In addition, swap data sets can be allocated to decrease swap times.

The importance of allocating swap data sets varies, depending on the level of the system and the nature of the configuration and workload. Except when you are using a fixed head device, these data sets should be isolated on separate volumes that are dedicated to ASM. To every degree possible, contention for the paths to these volumes should be minimized.

ASM logic arranges its data sets in circular queues ordered from those devices with the fastest service times (that is, cache and solid state device (SSD)) to those with the slowest (that is, fixed head then moveable head). When selecting a swap or page data set, the "fastest" queue is always scanned first. Other data sets, arrayed on the "slower" queues, are scanned only when no fast queue device can handle the request.

Within a queue, data sets are selected in round-robin sequence. In MVS/XA and MVS/ESA systems, an effort is made to load balance the paging subsystem by scheduling page-out requests to those devices that are projected to provide fast response times and bypassing those that are projected to provide slow response times.

ASM allocates available slots and builds paging channel programs in such a way that head movement is minimized. The generated channel programs contain SEEK and SET SECTOR commands as required.

When processing multiple paging requests, channel programs are constructed to run for a targeted period of time, called the service burst. Avoiding long-running channel programs allows ASM to spread the paging I/O workload, increase parallelism, and minimize forcing an otherwise innocent address space that is waiting for page fault resolution to wait for a long channel program to complete. The service burst concept is discussed in more detail in Section 2.3.4.2.2.

2.3.4.2 Paging

ASM defines three classes of paging space data sets, corresponding to the three logical groups of pages (pageable link pack area (PLPA), common, and local).

At least three paging data sets must be defined to the system. The first one is the PLPA data set, the second is the common data set, and the remaining data set(s) are used as local page data sets. You can define a fourth class of paging data set, the duplex data set, to provide a backup for the PLPA and common data set. Section 2.3.4 outlines what pages will be written to each of these data sets. In this section we discuss some device and data set size considerations.

The term slot is used to characterize the location of a virtual page when it resides on an ASM data set, just as the term frame is used when a virtual page resides in real memory. When a page is resident on an ASM data set, it is said to occupy a slot. ASM paging data sets are sized in units of slots.

PLPA DATA SET

The PLPA data set should be on a fixed-head device if possible. If not, make sure this data set is on a dedicated volume. Generally, a good starting point for sizing the PLPA data set is four megabytes. In MVS/XA systems, because of the potentially very large extended PLPA, this eventually may not be large enough, but it is a good place to start. If the PLPA data set is too small, it will spill into the common data set. Although this is not recommended, we give an example of how this might happen in Section 2.3.2.1.

COMMON DATA SET

The common data set should be large enough to hold all common system area and nucleus area pages (in MVS/XA the nucleus is conceptually included in the common area). As we noted for the PLPA data set, a reasonable starting point might be four megabytes (the same MVS/XA considerations apply), then adjust the size to the actual requirements.

When the common data set is too small, it will overflow to the PLPA data set. If this happens, ASM will not use the common data set again, even when slots become available. For this reason, you should avoid overflow from the common paging data set to the PLPA data set.

DUPLEX DATA SET

The duplex data set should be large enough to hold both PLPA and common pages. Thus, it should be as large as the combined size of the PLPA and common page data sets. Generally, eight megabytes should be sufficient for systems through MVS/ESA. Again, you can use the Auxiliary Storage Management reports to determine size requirements for this data set.

Use of duplexing is not generally recommended because of the extra overhead involved in the duplex operation (an extra page-out is done for each common area page-out). If you use duplexing, we recommend that you evaluate the savings in terms of system down time to see if this option is still required.

LOCAL DATA SET

The total paging space provided by the local data sets must be sufficient to hold all private area pages and virtual I/O (VIO) data sets. An example of this size calculation is given in the Initialization and Tuning Guides. You can use the Auxiliary Storage Management reports at the monthly level to refine this value.

The local data sets should be sized to reflect the type and number of address spaces expected to be active concurrently, the private area for each, and the total amount of VIO pages required. System and started task address spaces require local page data set slots equivalent to their private area size. To help ensure that there will be enough ASM slots (sometimes referred to as backing slots or backing store) available for the address space paging, swapping, and VIO demands, MVS performs the function described below.

In pre-MVS/XA systems, at address space creation for TS0 and batch, ASM divides the number of slots that would be required to contain the entire private area of the address space by a constant (ILRSLOTC, currently initialized to eight) and subtracts the result from the slot count for all local page data sets. The number of slots required for a VIO data set is divided by another constant (ILRSLOTV, currently initialized to four) and, again, the result is subtracted from the overall slot count. When there are no more slots available by this computation, further address spaces or VIO data sets are not allowed.

In MVS/XA and MVS/ESA systems, these functions are implemented through the ASMSLOT (for non-VIO requirements) and ASMSLOTV (for VIO requirements) values. The calculations are as described for pre-MVS/XA systems, but the default assumptions are different. It is assumed in both MVS/XA and MVS/ESA systems that you have allocated enough paging space to accommodate the needs of your workload. Hence the default values for both parameters effectively turn off the backing store functions. The method of turning them on and adjusting them to the needs of your workload are described in the Initialization and Tuning Guide.

This algorithm is NOT based on the actual number of local paging data set slots currently in use. Instead, it attempts to predict the number of slots that will be used by an address space or VIO data set and effectively reserves that number of slots.

The PLPA and common data sets should be sized somewhat above the minimum requirements that can be calculated. A growth factor can be determined from the ASM Page/Swap Activity Report and should be included. This growth factor is typically 25 percent to 35 percent of the maximum slots used over the last four months.

Sizing of local data sets is somewhat more complex. If the data set is going to service group requests (swapping, VIO, virtual fetch), it should be large enough to make efficient use of the consecutive slot allocation.

In sizing the PLPA and common data sets, you should be aware that overflow from one data set to the other is allowed. In general, only overflow from PLPA to common should be accepted.

The following sections discuss details of the paging algorithms used. Section 2.3.2 explains the application of these ideas.

- 1 - Data Set Selection
- 2 - Service Burst Calculation

2.3.4.2.1 Data Set Selection

For any page-in request, Auxiliary Storage Manager (ASM) must naturally select the page data set to which the virtual page was paged-out. To improve page fault response time, page-in requests are given priority over page-outs.

With the exception of a create LPA (CLPA) operation at IPL, pageable link pack area virtual pages are read-only, and are therefore paged-in from the PLPA data set, but never paged-out. Changed Common System Area (CSA) virtual pages are always paged-out to the common data set. The common and PLPA data sets may spill to one another, though common should not be allowed to spill to PLPA for performance reasons. For private area page-out processing, however, ASM must choose between the local data sets.

In MVS/XA and MVS/ESA, considerations are made for the projected response times of the individual page data sets. Page data sets with faster projected response times tend to be selected over those with slower projected response times, balancing the paging load to provide a more consistent level of paging response across page-out requests.

Naturally there also must be enough space on the selected page data set to service the request. For a group (or related) request, this means enough contiguous slots. For a demand paging request, there must be enough slots to hold the pages to be paged-out.

2.3.4.2.2 Service Burst Calculation

The idea behind the service burst is to avoid long-running paging channel programs while address spaces are waiting for page fault resolution. The goal is to limit paging channel program execution to a fixed number of pages, device dependent.

2.3.4.3 Swapping

Swapping is the means by which MVS frees real storage frames occupied by inactive or low priority address spaces and makes them available for use by higher priority active address spaces.

This section describes the swapping algorithms and explains swap data set configuration changes. A brief introduction to the evolution of the swap process provides a framework for this discussion. Some recommendations on swap data set sizing are given.

Before the availability of logical swapping and the extended swap IUP, only LSQA pages were written to swap data sets, and a physical swap-out was initiated when SRM analysis showed that an address space would be, or should be, made idle for an extended period. An address space was swapped:

- o Between transactions (for TSO address spaces).
- o For all explicit long wait conditions.
- o In other wait conditions that were "too long" (i.e., detected waits or ENQ contention).
- o Unilaterally by the SRM for multiprocessing level contention and resource under or over utilization.

A swap-out is performed in three steps:

1. Execution is halted if the address space is not already in a wait state and the address space is quiesced (i.e., put in a wait state).
2. The pages held by the address space are trimmed.
Non-LSQA pages unchanged since their last page-in have their frames placed on the available frame queue, where they are immediately available for stealing.
3. Non-LSQA working set pages that have been changed since their last page-in are paged-out to the local data sets, and their slot locations recorded in a control block in the LSQA. The LSQA is then paged-out to the swap data sets, if available; otherwise, the LSQA is paged-out to the local data sets.

The swap-in process is the reverse of a swap-out:

1. LSQA is paged-in.
2. Page-ins are scheduled for non-LSQA pages marked as referenced at swap-out.
3. The address space is restored and marked dispatchable, possibly before the page-ins from step two have completed.

Starting with MVS/SE2, logical swapping was introduced for TSO address spaces. The effect of this change was to avoid physically swapping a TSO address space if: (1) the user is actively entering commands, and (2) there is enough real storage available.

A second major change was in the form of an installed user program (IUP) to modify the swap process. With the extended swap IUP, those pages marked as referenced at swap-out are treated as LSQA and paged-out to the swap data sets. Swap-in step one brings in all the pages necessary to continue execution of the address space, so swap-in step two is effectively skipped. Swap-in step three is also significantly shortened, because it is unnecessary to wait for page-ins scheduled by step two.

The swap process is discussed in the following sections.

- 1 - Physical Swap Algorithm
- 2 - Working Set Trim
- 3 - Logical Swapping

2.3.4.3.1 Physical Swap Algorithm

Ignoring the quiesce/restore and trimming steps, the swap process of the extended swap IUP and MVS/SP 1.3 and beyond is really a two-stage swap-out and a single-stage swap-in.

Changed but unreferenced pages are paged-out to local data sets, saving the ASM slot locations in LSQA control blocks. MVS/XA and MVS/ESA use the sequential slot allocation algorithm to select slots on the local page data sets.

Then LSQA and non-LSQA referenced pages are paged-out to swap data sets. All the pages needed for a swap-in receive the advantage of the sequential slots of swap sets.

If multiple swap sets are required for the larger working sets, they will be split among multiple swap data sets and the I/O will run concurrently (assuming a sufficient number of allocated data sets). When an address space is swapped back into real memory, all of its pages are brought in from the swap data sets in a single stage operation.

2.3.4.3.2 Working Set Trim

In this section the term working set is used to be compatible with the IBM references on ASM operation. You should be aware that this is a third definition of working set. This version of the working set is a pragmatic attempt to keep enough pages at swap-out to ensure minimal paging at swap-in. As such, the swap working set determination includes more pages than might be in the Denning working set.

Working set trim makes page frames available for immediate stealing and indirectly determines the number of non-LSQA pages that will be swapped. An address space can be trimmed down to its minimum for storage isolation (see the PWSS parameter in the Initialization and Tuning Guide) if such a minimum was specified.

Working set trim attempts to selectively protect unreferenced pages from being stolen. Referenced pages are never stolen, because they must be paged-out to the swap data sets by the extended swap algorithm. The method by which working set trim discriminates between pages that should and should not be made available for stealing is different for physical and logical swaps.

PHYSICAL SWAP WORKING SET TRIM

For address spaces that are being physically swapped-out, working set trim is based on elapsed time since the last unreferenced interval count (UIC) update. When a physical swap is initiated within 1/2 SRM second of the last UIC update, then unreferenced pages with a UIC of zero are not trimmed. This protects pages that might be frequently referenced, but were unlucky enough to have their reference bit turned off by the very recent UIC update. It is assumed there has not been sufficient time for the page to be referenced again.

LOGICAL SWAP WORKING SET TRIM

For address spaces that are being logically swapped-out, working set trim is based on current utilization of real storage. Unreferenced pages within a UIC of up to 23 are not trimmed if their UIC is less than the integer part of the following calculation:

$$1 + (\text{System Average Max UIC} - 30) / 10$$

Thus, logically swapped address spaces are always trimmed

less severely than physically swapped address spaces because this calculation has a minimum value of one. The purpose of this algorithm is to protect more pages from being trimmed as more real storage is made available.

2.3.4.3.3 Logical Swapping

MVS/SE1 introduced the concept of logical swapping. The consumption of CPU and I/O resources by swapping had limited the throughput of large TSO systems, primarily because completion of each TSO transaction initiated a physical swap-out. The purpose of logical swapping is, where possible, to avoid actual swap-out page transfer to the ASM data sets. To describe logical swapping, we first consider the steps in the life of a transaction in a system prior to MVS/SE1:

1. Swap-in of the LSQA pages (stage 1).
2. Swap-in of the private pages (stage 2).
3. Restore processing - ready the address space for execution.
4. Execution phase.
5. Quiesce processing - stop all tasks, SRBs, and I/O.
6. Working set trim.
7. Swap-out of changed non-LSQA pages and LSQA pages.

The decision to do a physical or logical swap is made at the completion of quiesce processing and is based on two variables: (1) the user think time and (2) the system think time. Prior to MVS/SP 1.3, only TSO address spaces that were quiesced due to a terminal wait condition (input or output) were considered for logical swapping.

Starting in MVS/SP 1.3, any swappable task that enters a detected or long wait condition is considered for logical swapping. User think time is calculated as the time from the end of quiesce processing to the end of restore processing, and hence accounts for user idle time, time keying in data, swap-in time, and any communications delays.

System think time is calculated as a function of the available storage. Consult the Initialization and Tuning Guide for the appropriate level of your system for the exact calculation used. The value of system think time may be plotted over time. This value is controlled by parameters in the SYS1.PARMLIB member IEAOPTnn that define the range of values, the length of the available frame queue, and the system unreferenced interval count (UIC) value.

With MVS/SP 1.3 and beyond, the amount of fixed storage and the amount of storage fixed or allocated from page-in or page-out operations is also considered. An address space is logically swapped if the previous user think time is less than the current value of system think time. When a logically swapped address space remains idle for longer than the system think time, it is physically swapped.

TSO users receive an extra 2-second grace period, but the physical swap is not actually performed until real storage frames are needed. Hence, an address space may remain logically swapped far longer than might be expected.

Non-TSO address spaces are also candidates for logical swapping. The criteria for non-TSO logical swapping are:

- o System think time must be greater than five seconds.
- o The current high UIC must be greater than logical swap high threshold (default is 30).
- o The average number of fixed frames below 16 megabytes is below that logical swap threshold.

When all the above conditions are met, a swappable non-TSO address space entering a detected or long wait will be considered for logical swapping.

Given sufficient real storage, logical swapping greatly improves TSO response. There are so many interactions between an ASM configuration and the amount of real memory, the type of swapping (extended swap or not), and page stealing that it is difficult to give firm guidelines on how to use these facilities. You should perform all basic system tuning first, however, before you make any changes to the system-provided defaults. Some general strategies are discussed in Section 2.3.2 of this guide.

In an expanded storage environment, logical swapping is still utilized in an attempt to save the processor cycles and expanded storage or I/O overhead of physical swapping. As noted above, a TSO address space in a terminal wait condition or any class of address space in a long or detected wait condition is eligible for logical swapping. When this happens, the LSQA, fixed, and most recently referenced pages are kept in central storage. The only real impact of expanded storage on the logical swap process is to mitigate the degradation that might occur when a logical swap must be

converted to a physical swap.

Just as a logical swap is effective only if the address space is swapped back in before being converted to a physical swap, an expanded storage physical swap is effective only if the address space is swapped back in before being migrated through central to auxiliary storage.

2.3.5 Case Study

The following case study details the application of the Auxiliary Storage Management routines to review changes made to an ASM configuration.

These topics are discussed in detail:

- 1 - Problem Description
- 2 - Inquiry Composition
- 3 - Output Reports

2.3.5.1 Problem Description

The purpose of this study is to review the ASM configuration for a 3033 running MVS/SP 1.1 with the extended swap IUP. The configuration was modified to remove two drums and two 3350 swap data sets. The new ASM configuration is isolated on channel 1 with an alternate path on channel 7. It was decided to place the common and PLPA data sets on the same volume. Because this system is at an SP 1.1 level, this is appropriate (see Section 2.3.4.1 for SP 1.3 considerations).

The study has the following objectives:

- o Ensure that local paging does not suffer from the removal of the drums.
- o Look for page response problems.
- o Verify that "low activity" data sets on the swap volumes have not caused problems.
- o Check for any overloading conditions.
- o Review the initial size estimates.

The first step is to use the ASM Page/Swap Response Time Report to produce a before and after change overview of the configurations. We chose the WEEKS time-span because gross trending is our basic objective. To review the operation and service given by each of the individual data sets, we produce the ASM Page/Swap Activity Report at an hourly level.

2.3.5.2 Inquiry Composition

Figures 2-72 through 2-116 detail the options that we provided to the Auxiliary Storage Management inquiry panels. These options are discussed in the paragraphs below.

The SYSID selection parameter limits the reports to the system under study. We selected Zone 1 (prime time) because this system runs TSO and some online DBMS applications during the day and mostly production batch at night. Thus, the heavy paging and swapping load comes during prime time.

The entries in the ASM Page/Swap Response Time Report inquiry panel in Figure 2-72 produce the report shown in Figure 2-25, which summarizes the old configuration. Each cycle in the WEEKS timespan is one week, thus a report covering 10-14 weeks ago is produced by these statements.

```

/----- ASM Page/Swap Response Time Report -----\
| Command ==>
|
| Composing CA MICS Inquiry:  ASMC1  - Response Time / Old Configuration
|
| Time-span ==> WEEKS  (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 10 - 14
| Data base ==> A (PRIMARY)  _          _          _
| Allow Run-Time Override of Data Base ==> Y (Y/N) and/or Cycle(s) ==> Y (Y/N)
|
| Report by hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone      ==> Y (Y/N)
|
| Report level (Select only one):  Page/Swap data set names ==> Y (Y/N)
|                                Page/Swap data set types  ==> N (Y/N)
|                                Page/Swap device types    ==> N (Y/N)
|
| Data Selection:                                Allow run-time override
| Date          ==> _____ - _____ - _____ ==> N (Y/N)
| Hour          ==> __ - __ - __ - __ - __ ==> N (Y/N)
| Zone          ==> 1 _ _ _ _ _ ==> N (Y/N)
| Sysid         ==> PROD _____ ==> N (Y/N)
|
| Specify extended options ==> N (Y/N)
\-----/

```

Figure 2-72. Case Study - ASM Page/Swap Response Time Report Inquiry Panel / Old Configuration

For a full description of the fields on this panel, see Section 2.3.6, Component Operation, or press the help key (usually PF1) when you are at this panel.

The entries in the ASM Page/Swap Response Time Report inquiry panel in Figure 2-73 produce the report shown in Figure 2-76 covering the current week and the previous two weeks for ZONE 1.

```

/----- ASM Page/Swap Response Time Report -----\
| Command ==>                                     |
| Composing CA MICS Inquiry:  ASMC2  - Response Time / New Configuration |
| Time-span ==> WEEKS  (DETAIL/DAYS/WEEKS/MONTHS)      Cycle(s) ==> 01 - 03 |
| Data base ==> A (PRIMARY)  _          _          _          _          |
| Allow Run-Time Override of Data Base ==> Y (Y/N) and/or Cycle(s) ==> Y (Y/N) |
| Report by hour if timespan is DAYS or WEEKS ==> N (Y/N) |
| Generate a report for each value of Zone      ==> Y (Y/N) |
| Report level (Select only one):  Page/Swap data set names  ==> Y (Y/N) |
|                                   Page/Swap data set types   ==> N (Y/N) |
|                                   Page/Swap device types      ==> N (Y/N) |
| Data Selection:                                     Allow run-time override |
| Date           ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
| Hour           ==> _ _ - _ _ - _ _ - _ _ - _ _ - _ _ ==> N (Y/N) |
| Zone           ==> 1 _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
| Sysid          ==> PROD _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-73. Case Study - ASM Page/Swap Response Time Report Inquiry Panel / New Configuration

The entries in the ASM Page/Swap Activity Report inquiry panel in Figure 2-74 produce a report for each data set in the ASM configuration. These reports are shown in Figure 2-77.

```

/----- ASM Page/Swap Activity Report -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: ASMC3 - ASM Activity by Data Set |
| Time-span ==> DAYS (DETAIL/DAYS/WEEKS/MONTHS)      Cycle(s) ==> 01 - 01 |
| Data base ==> A (PRIMARY) _                      _ |
| Allow Run-Time Override of Data Base ==> Y (Y/N) and/or Cycle(s) ==> Y (Y/N) |
| Report by hour if timespan is DAYS or WEEKS ==> Y (Y/N) |
| Generate a report for each value of Zone ==> Y (Y/N) |
| Report level (Select only one): Page/Swap data set names ==> Y (Y/N) |
|                               Page/Swap data set types ==> N (Y/N) |
|                               Page/Swap device types ==> N (Y/N) |
| Data Selection:                                     Allow run-time override |
| Date ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
| Hour ==> _ _ - _ _ - _ _ - _ _ - _ _ ==> N (Y/N) |
| Zone ==> 1 _ _ _ _ _ ==> N (Y/N) |
| Sysid ==> PROD _ _ _ _ _ ==> N (Y/N) |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-74. Case Study - ASM Page/Swap Activity Report Inquiry Panel / By Data Set

2.3.5.3 Output Reports

Figures 2-75 and 2-76 show the output from the two ASM Page/Swap Response Time Reports. From these figures, we can see that the change to the final configuration was made in two steps. The drums were removed during the week of 9/7/yy and then, by the week of 11/15/yy, the ASM configuration was isolated on channel 1 to reach the final configuration shown in the second ASM Page/Swap Response Time Report (Figure 2-76).

As we noted at the beginning of this case study, in setting up the new configuration we had the following goals:

- o Ensure that local paging does not suffer from removal of the drums.
- o Look for page response problems.
- o Verify that the "low activity" data sets on the swap volumes have not caused problems.
- o Check for any overloading conditions.
- o Review the initial size estimates.

First, the 3380s are giving approximately the same service as the drum, so clearly local paging was not hurt by this change. We should review hourly ASM Page/Swap Response Time Reports to ensure that consistent service is being given. Next, based on the data set busy times and page transfer rates, this configuration is rather lightly loaded.

As a final note, by isolating the ASM data sets on a channel, we should get more consistent response.

CA MICS Performance Manager WEEKLY ASM DATASET RESPONSE REPORT FOR ZONE = 1 SYSTEM = PROD RUN DATE = NOV 20, yyyy						
WEEKS:	08/18/yy	08/25/yy	09/01/yy	09/08/yy	09/15/yy	CHANGE %
COMMON DATASET						
SYS2.LOCAL181 (3380)	0.017	0.023	0.021	0.017	0.023	35.29 %
LOCAL DATASET						
SYS2.LCL1A0A (2305-2)	0.030	0.038	0.047	0.000	0.000	N/A
SYS2.LCL1A0B (2305-2)	0.029	0.038	0.041	0.000	0.000	N/A
SYS2.LOCAL168 (3380)	0.000	0.000	0.023	0.020	0.021	-8.70 %
SYS2.LOCAL182 (3380)	0.029	0.032	0.034	0.026	0.026	-10.34 %
SYS2.LOCAL183 (3380)	0.030	0.033	0.034	0.032	0.029	-3.33 %
PLPA DATASET						
SYS2.PLPA180 (3380)	0.016	0.022	0.026	0.019	0.018	12.50 %
SWAP DATASET						
SYS2.SWAP16C (3380)	0.000	0.000	0.075	0.067	0.067	-10.60 %
SYS2.SWAP291 (3380)	0.035	0.035	0.040	0.033	0.037	5.71 %
SYS2.SWAP292 (3380)	0.043	0.045	0.049	0.041	0.041	-4.65 %
SYS2.SWAP5C3 (3350)	0.108	0.113	0.127	0.000	0.000	N/A
SYS2.SWAP3EA (3350)	0.077	0.083	0.081	0.106	0.077	---

Figure 2-75. Case Study - ASM Page/Swap Response Time Report / Old Configuration

CA MICS Performance Manager				
WEEKLY ASM DATASET RESPONSE REPORT				
FOR ZONE = 1				
SYSTEM = PROD		RUN DATE = NOV 20, yyyy		

WEEKS:	11/03/yy	11/10/yy	11/17/yy	CHANGE %

COMMON DATASET				

SYS2.LOCAL161 (3380)	0.021	0.021	0.019	-9.52 %

LOCAL DATASET				

SYS2.LOCAL16C (3380)	0.021	0.022	0.019	-9.52 %
SYS2.LOCAL16E (3380)	0.020	0.021	0.021	5.00 %
SYS2.LOCAL168 (3380)	0.021	0.020	0.019	-9.52 %

PLPA DATASET				

SYS2.PLPA161 (3380)	0.025	0.027	0.039	56.00 %

SWAP DATASET				

SYS2.SWAP17D (3380)	0.060	0.059	0.060	0.00 %
SYS2.SWAP171 (3380)	0.055	0.051	0.056	1.82 %
SYS2.SWAP173 (3380)	0.125	0.114	0.117	-6.37 %

Figure 2-76. Case Study - ASM Page/Swap Response Time Report / New Configuration

CA MICS Performance Manager											
HOURLY ASM DATASET ACTIVITY REPORT											
FOR ZONE = 1											
SYSTEM = PROD						RUN DATE = NOV 20, yyyy					
DSN = SYS1.PAGE16C UNIT = 16C						DEVICE TYPE = 3380 PAGE SPACE TYPE = LOCAL					

HOUR/DATE (HR-MM/DD): 08-11/19 09-11/19 10-11/19 11-11/19 12-11/19 13-11/19 14-11/19 15-11/19 16-11/19 17-11/19 CHANGE %											

DATASET SIZE											

AVG SLOTS ALLOCATED	12750	12750	12750	12750	12750	12750	12750	12750	12750	12750	---
AVG SLOTS USED	2086.25	2709.00	2950.75	3477.00	3368.50	3475.75	3498.25	3637.00	3497.25	3071.63	47.24 %
MAX SLOTS USED	2612.00	2898.00	3278.00	3858.00	3695.00	3789.00	3716.00	3857.00	3879.00	3257.00	24.69 %
UNUSABLE SLOTS	0	0	0	0	0	0	0	0	0	0	N/A
PERCENT SLOTS USED	16.36 %	21.25 %	23.14 %	27.27 %	26.42 %	26.26 %	27.44 %	28.53 %	27.43 %	24.09 %	47.24 %

DATASET WORKLOAD											

TOTAL PAGING I/O	2430	5401	9494	13162	11164	16030	17493	18439	14192	7795	220.78 %
TOTAL DEVICE I/O	2433	5452	9498	13240	11166	16052	17507	18442	14194	7820	221.41 %
I/O CONTENTION INDEX	0.001	0.009	0.000	0.006	0.000	0.001	0.001	0.000	0.000	0.003	300.00 %
AVG I/O RATE	0.68	1.50	2.64	3.66	3.10	4.45	4.86	5.12	3.94	2.89	327.71 %
MAX I/O RATE	0.99	1.77	3.98	4.51	3.66	5.63	5.98	6.47	5.41	3.30	232.27 %
AVG PAGE RATE	3.60	2.52	0.37	2.13	3.42	2.41	3.11	2.44	1.54	2.68	-25.55 %
AVG PAGE BURST SIZE	12.81	13.40	14.00	14.78	14.51	15.88	14.88	13.21	11.65	12.71	-0.01 %
AVG DATASET UTIL	1.50 %	2.33 %	5.58 %	7.86 %	5.94 %	10.69 %	9.47 %	10.88 %	8.13 %	4.96 %	230.86 %
MAX DATASET UTIL	3.00 %	3.44 %	7.65 %	10.44 %	7.66 %	12.45 %	12.22 %	13.76 %	11.76 %	5.77 %	92.59 %

DATASET SERVICE											

AVG RESPONSE TIME/I/O	0.023	0.016	0.022	0.022	0.019	0.024	0.020	0.021	0.021	0.018	-21.74 %
AVG SERVICE TIME/I/O	0.022	0.016	0.021	0.022	0.019	0.024	0.019	0.021	0.021	0.017	-22.64 %
AVG SERVICE TIME/PAGE	0.002	0.001	0.002	0.004	0.003	0.003	0.002	0.003	0.002	0.001	-50.00 %
MAX SERVICE TIME	0.030	0.020	0.026	0.023	0.021	0.029	0.021	0.022	0.024	0.020	-35.06 %

Figure 2-77. Case Study - ASM Page/Swap Data Set Activity Report (Part 1 of 4)

CA MICS Performance Manager											
HOURLY ASM DATASET ACTIVITY REPORT											
FOR ZONE = 1											
SYSTEM = PROD				RUN DATE = NOV 20, yyyy							
DSN = SYS1.PAGE161				UNIT = 161	DEVICE TYPE = 3380		PAGE SPACE TYPE = COMMON				

HOUR/DATE (HR-MM/DD): 08-11/19 09-11/19 10-11/19 11-11/19 12-11/19 13-11/19 14-11/19 15-11/19 16-11/19 17-11/19 CHANGE %											

DATASET SIZE											

AVG SLOTS ALLOCATED	12750	12750	12750	12750	12750	12750	12750	12750	12750	12750	---
AVG SLOTS USED	1648.50	1655.50	1667.00	1695.25	1707.25	1710.00	1717.75	1729.44	1729.75	1662.00	0.82 %
MAX SLOTS USED	1654.00	1660.00	1674.00	1709.00	1710.00	1715.00	1728.00	1732.00	1732.00	1728.00	4.47 %
UNUSABLE SLOTS	0	0	0	0	0	0	0	0	0	0	N/A
PERCENT SLOTS USED	12.93 %	12.98 %	13.07 %	13.30 %	13.39 %	13.41 %	13.47 %	13.56 %	13.57 %	13.04 %	0.82 %

DATASET WORKLOAD											

TOTAL PAGING I/O	467	1094	3091	6675	4214	7753	8286	9919	6155	2014	331.26 %
TOTAL DEVICE I/O	475	1098	3092	6695	4219	7766	8487	9931	6156	2085	338.95 %
I/O CONTENTION INDEX	0.017	0.004	0.000	0.003	0.001	0.002	0.024	0.001	0.000	0.035	105.88 %
AVG I/O RATE	0.13	0.30	0.86	1.85	1.17	2.15	2.30	2.76	1.71	0.75	475.09 %
MAX I/O RATE	0.20	0.46	1.41	2.90	1.46	3.15	3.38	3.58	2.57	1.07	423.14 %
AVG PAGE RATE	3.39	2.64	1.51	0.73	0.34	0.31	0.45	0.56	0.87	0.66	-80.53 %
AVG PAGE BURST SIZE	19.54	15.32	16.78	14.33	16.78	14.79	14.48	15.61	14.08	12.83	-34.33 %
AVG DATASET UTIL	0.30 %	0.55 %	1.58 %	3.83 %	2.08 %	3.85 %	4.91 %	5.50 %	3.38 %	1.40 %	360.57 %
MAX DATASET UTIL	0.77 %	1.00 %	2.33 %	6.33 %	2.66 %	5.77 %	8.00 %	6.99 %	5.33 %	1.55 %	99.98 %

DATASET SERVICE											

AVG RESPONSE TIME/I/O	0.024	0.019	0.020	0.021	0.018	0.019	0.023	0.021	0.020	0.019	-19.87 %
AVG SERVICE TIME/I/O	0.024	0.018	0.018	0.021	0.018	0.018	0.021	0.020	0.020	0.019	-19.90 %
AVG SERVICE TIME/PAGE	0.001	0.001	0.001	0.002	0.001	0.001	0.002	0.001	0.001	0.001	0.00 %
MAX SERVICE TIME	0.038	0.029	0.026	0.022	0.024	0.020	0.025	0.022	0.023	0.026	-31.56 %

Figure 2-77. Case Study - ASM Page/Swap Data Set Activity Report (Part 2 of 4)

CA MICS Performance Manager											
HOURLY ASM DATASET ACTIVITY REPORT											
FOR ZONE = 1											
SYSTEM = PROD				RUN DATE = NOV 20, yyyy							
DSN = SYS1.PAGE168				UNIT = 168	DEVICE TYPE = 3380	PAGE SPACE TYPE = LOCAL					

HOUR/DATE (HR-MM/DD): 08-11/19 09-11/19 10-11/19 11-11/19 12-11/19 13-11/19 14-11/19 15-11/19 16-11/19 17-11/19 CHANGE %											

DATASET SIZE											

AVG SLOTS ALLOCATED	12750	12750	12750	12750	12750	12750	12750	12750	12750	12750	---
AVG SLOTS USED	2156.75	2793.50	3107.50	3567.50	3437.00	3503.50	3546.50	3693.94	3696.19	3264.31	51.35 %
MAX SLOTS USED	2456.00	3004.00	3391.00	3902.00	3688.00	3725.00	3900.00	3964.00	3966.00	3496.00	42.35 %
UNUSABLE SLOTS	0	0	0	0	0	0	0	0	0	0	N/A
PERCENT SLOTS USED	16.92 %	21.91 %	24.37 %	27.98 %	26.96 %	27.48 %	27.82 %	28.97 %	28.99 %	25.59 %	51.35 %

DATASET WORKLOAD											

TOTAL PAGING I/O	2544	5423	9994	13884	11634	17101	17061	18197	14451	8043	216.16 %
TOTAL DEVICE I/O	2552	5426	10056	13891	11636	17103	17062	18197	14465	8065	216.03 %
I/O CONTENTION INDEX	0.003	0.001	0.006	0.001	0.000	0.000	0.000	0.000	0.001	0.005	66.67 %
AVG I/O RATE	0.71	1.51	2.78	3.86	3.23	4.75	4.74	5.05	4.01	2.98	321.54 %
MAX I/O RATE	0.93	1.83	3.96	4.77	3.91	5.82	5.88	6.24	5.13	3.18	241.51 %
AVG PAGE RATE	2.39	2.74	1.43	0.89	1.23	0.61	0.55	0.57	1.87	1.66	-30.54 %
AVG PAGE BURST SIZE	17.54	15.76	14.76	14.54	15.66	13.65	14.33	12.55	13.08	14.73	-16.02 %
AVG DATASET UTIL	1.50 %	2.83 %	5.69 %	7.91 %	6.86 %	9.66 %	9.86 %	10.19 %	9.05 %	6.11 %	307.41 %
MAX DATASET UTIL	2.11 %	3.66 %	8.87 %	10.44 %	8.22 %	12.90 %	12.55 %	12.22 %	10.98 %	6.55 %	210.52 %

DATASET SERVICE											

AVG RESPONSE TIME/I/O	0.022	0.019	0.021	0.021	0.022	0.021	0.021	0.021	0.023	0.021	-3.36 %
AVG SERVICE TIME/I/O	0.021	0.019	0.021	0.021	0.021	0.020	0.021	0.020	0.023	0.021	-2.23 %
AVG SERVICE TIME/PAGE	0.001	0.001	0.001	0.001	0.001	0.002	0.001	0.002	0.002	0.001	0.00 %
MAX SERVICE TIME	0.024	0.022	0.023	0.023	0.023	0.022	0.022	0.023	0.026	0.021	-11.14 %

Figure 2-77. Case Study - ASM Page/Swap Data Set Activity Report (Part 3 of 4)

CA MICS Performance Manager											
HOURLY ASM DATASET ACTIVITY REPORT											
FOR ZONE = 1											
SYSTEM = PROD				RUN DATE = NOV 20, yyyy							
DSN = SYS1.SWAP17D				UNIT = 17D	DEVICE TYPE = 3380	PAGE SPACE TYPE = SWAP					

HOUR/DATE (HR-MM/DD): 08-11/19 09-11/19 10-11/19 11-11/19 12-11/19 13-11/19 14-11/19 15-11/19 16-11/19 17-11/19 CHANGE %											

DATASET SIZE											

AVG SLOTS ALLOCATED	240	240	240	240	240	240	240	240	240	240	---
AVG SLOTS USED	46.00	57.50	63.25	66.50	59.50	62.25	61.25	62.00	69.75	58.33	26.79 %
MAX SLOTS USED	57.00	69.00	78.00	83.00	72.00	75.00	81.00	78.00	87.00	74.00	29.82 %
UNUSABLE SLOTS	0	0	0	0	0	0	0	0	0	0	N/A
PERCENT SLOTS USED	19.17 %	23.96 %	26.35 %	27.71 %	24.79 %	25.94 %	25.52 %	25.83 %	29.06 %	24.30 %	26.79 %

DATASET WORKLOAD											

TOTAL PAGING I/O	1424	2617	3442	4073	3507	4743	4987	5004	4960	3028	112.64 %
TOTAL DEVICE I/O	1435	2617	3449	4105	3515	4748	5012	5005	4963	3115	117.07 %
I/O CONTENTION INDEX	0.008	0.000	0.002	0.008	0.002	0.001	0.005	0.000	0.001	0.029	262.50 %
AVG I/O RATE	0.40	0.73	0.96	1.13	0.97	1.32	1.39	1.39	1.38	1.12	183.47 %
MAX I/O RATE	0.52	0.88	1.04	1.32	1.04	1.50	1.55	1.55	1.57	1.21	131.42 %
AVG PAGE RATE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AVG PAGE BURST SIZE	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AVG DATASET UTIL	1.38 %	2.75 %	3.33 %	3.75 %	3.47 %	5.19 %	5.97 %	5.44 %	5.27 %	4.51 %	225.37 %
MAX DATASET UTIL	2.44 %	2.88 %	4.33 %	4.33 %	4.22 %	5.99 %	7.11 %	6.65 %	6.00 %	5.11 %	109.41 %

DATASET SERVICE											

AVG RESPONSE TIME/I/O	0.051	0.052	0.053	0.048	0.057	0.060	0.066	0.060	0.059	0.063	23.53 %
AVG SERVICE TIME/I/O	0.043	0.045	0.043	0.040	0.046	0.050	0.054	0.049	0.048	0.051	18.62 %
AVG SERVICE TIME/PAGE	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	N/A
MAX SERVICE TIME	0.069	0.054	0.059	0.045	0.056	0.062	0.060	0.059	0.052	0.058	-15.96 %

Figure 2-77. Case Study - ASM Page/Swap Data Set Activity Report (Part 4 of 4)

2.3.6 Component Operation

This Component Operation Section describes the online inquiry panels required for the operation of Auxiliary Storage Management Analysis routines.

Entries in the online inquiry panels define the types of reports to be generated. In particular, these panels define:

- o The cycles from the CA MICS SCPPSD and HARDVA Files to be processed.
- o The summarization level of the report as hourly, daily, weekly, or monthly. If you have selected the DETAIL timespan, the data is summarized based on the length of your RMF interval. Note that the availability of each summarization level assumes that your system supports the required files at each level.
- o The type of report to be produced.

In addition to the timespan qualifier, a summarization option is provided in the form of ZONE and HOUR specification. You can specify summarization to the ZONE level for all timespans, and to the HOUR level for the WEEKS or DAYS timespan, with the Report by Hour and the Generate a report for each value of Zone options.

If you specify the DAYS timespan and HOUR summarization, and you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized to the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

The Auxiliary Storage Management reports provided are:

- o ASM Page/Swap Activity Report - Describes the activity and response for each data set found in the SCPPSD File.
- o ASM Page/Swap Response Time Report - Reports the device response of each data set found in the SCPPSD File.

The following options are common to both control panels:

Timespan: The timespan (DETAIL, DAYS, WEEKS, or MONTHS) of the CA MICS file(s) that you wish to use as

input for processing.

Cycle(s):	One or more cycles of the CA MICS files that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a cycle from the CA MICS database will usually contain a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided by the DATE and HOUR options.
Database:	The single-character CA MICS identifier for the database to be accessed. Valid values are the database unit identifiers available at your installation. Contact your CA MICS System Administrator if you do not know what database units are available, or enter a question mark (?) to view a list of your database units. You can select up to five databases per inquiry.
Allow Run-Time Override:	The indication of whether the database identifier and/or the Cycles specification can be overridden at inquiry execution time.
Report by hour:	The indication of whether a report should be generated for each hour. This is only valid for the DAYS or WEEKS timespans.
Generate a report for each value of Zone:	The indication of whether an individual report should be generated for each zone found in the input data or whether all zones should be merged together for the report.
Date:	Up to two date ranges (DDMONYY format) to be used for analysis. This option is normally used to drop the extra day from a cycle that is due to the timing of the SMF dump.
Hour:	One to three ranges of hours for which data is to be reported. For an individual hour, leave the end hour blank or make it the same as the start hour. You can specify a value from 00 to 23.
Zone:	The specification that zone selection is to be applied. If zone 1 definition is 8:00 a.m. to 6:00 p.m., you can specify reports covering that time period by entering a 1 in the ZONE field.

The following sections describe the control panels for each ASM report:

- 1 - ASM Page/Swap Activity Report
- 2 - ASM Page/Swap Response Time Report

2.3.6.1 ASM Page/Swap Activity Report

The ASM Page/Swap Activity Report describes the activity and response of each data set found in the SCPPSD Files.

The ASM Page/Swap Activity Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options that are normally defaulted.

```

/----- ASM Page/Swap Activity Report -----
| Command ==>
|
| Composing CA MICS Inquiry: ASMT1 - Run Activity Report
|
| Time-span ==> DETAIL (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 02
| Data base ==> A (PRIMARY)      _      _      _      _
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N)
|
| Report by hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone      ==> Y (Y/N)
|
| Report level (Select only one):   Page/Swap data set names ==> Y (Y/N)
|                                   Page/Swap data set types  ==> N (Y/N)
|                                   Page/Swap device types     ==> N (Y/N)
|
| Data Selection:                                Allow run-time override
|   Date      ==> _____ - _____ - _____ ==> N (Y/N)
|   Hour      ==> __ - __      __ - __      __ - __      ==> N (Y/N)
|   Zone      ==> _ - _ - _ - _ - _ - _ - _ - _ - _ - _ ==> N (Y/N)
|   Sysid     ==> _____ - _____ - _____ ==> N (Y/N)
|
| Specify extended options ==> N (Y/N)
\-----

```

Figure 2-79. ASM Page/Swap Activity Report Inquiry Panel

Figure 2-79 shows the ASM Page/Swap Activity Report panel. The report level options you can select for this report are as follows:

Report by data set names:	A report is to be generated for each data set found in the page/swap data set (SCPPSD) files.

Report by The page/swap data set information is to be

data set types:	summed by data set type: LOCAL, COMMON, PLPA, DUPLEX, and SWAP. This option allows the operation of the ASM system to be viewed as a single unit (e.g., what is the response for local paging?).
Report by device types:	The page/swap dataset information is to be summed by device type. This option is intended to aid you in evaluating the potential benefits of using different device types (e.g., replacing 2305-2 with 3380s for local data sets).

There are four other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.3.6.2 ASM Page/Swap Response Time Report

The ASM Page/Swap Response Time Report describes the response time for each data set found in the SCPPSD Files.

The ASM Page/Swap Response Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options that are normally defaulted.

```

/----- ASM Page/Swap Response Time Report -----
| Command ==>
|
| Composing CA MICS Inquiry: ASMT2 - Response Time Report
|
| Time-span ==> DETAIL (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 02
| Data base ==> A (PRIMARY)      _      _      _      _
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N)
|
| Report by hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone      ==> Y (Y/N)
|
| Report level (Select only one):   Page/Swap data set names ==> Y (Y/N)
|                                   Page/Swap data set types  ==> N (Y/N)
|                                   Page/Swap device types     ==> N (Y/N)
|
| Data Selection:                                Allow run-time override
|   Date      ==> _____ - _____ - _____ ==> N (Y/N)
|   Hour      ==> __ - __      __ - __      __ - __      ==> N (Y/N)
|   Zone      ==> _ - _ - _ - _ - _ - _ - _ - _ - _ - _ ==> N (Y/N)
|   Sysid     ==> _____ - _____ - _____ ==> N (Y/N)
|
| Specify extended options ==> N (Y/N)
\-----/

```

Figure 2-80. ASM Page/Swap Response Time Report Inquiry Panel

Figure 2-80 shows the ASM Page/Swap Response Time Report inquiry panel. The fields common to all ASM Analysis routines are described in Section 2.3.6. The report level options you can select for this report are as follows:

Report by data set names:	A report is to be generated for each data set found in the page/swap data set (SCPPSD) files.

Report by data set	The page/swap data set information is to be summed by data set type: LOCAL, COMMON, PLPA,

types: DUPLEX, and SWAP. This option allows the operation of the ASM system to be viewed as a single unit (e.g., what is the response for local paging?).

Report by device types: The page/swap data set information is to be summed by device type. This option is intended to aid you in evaluating the potential benefits of using different device types (e.g., replacing 2305-2 with 3380s for local data sets).

There are four other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.3.7 Requirements

The Auxiliary Storage Management reports give you the capability to track the allocation of auxiliary storage on a monthly, weekly, daily, hourly, or interval basis.

The reports use data from the page/swap data set files and from the device activity files from the RMF SCP and HAR Information Areas to report on the activity levels and response of the ASM subsystem. The device activity files are also used to add in the I/O queue time for each of the page/swap devices, because the I/O service time in the page/swap files includes only the I/O transfer time. Figure 2-81 shows the CA MICS files and timespan levels that are required to produce the ASM Analysis reports.

CA MICS FILE	LEVEL				
	X	D	W	M	Y
HARDVA	x	x	x	x	
SCPPSD	x	x	x	x	

Figure 2-81. ASM Analysis CA MICS File Requirements

For the CA MICS HARDVA and HARPSD Files, only one level is required (i.e., DETAIL, DAYS, WEEKS, or MONTHS), depending on the timespan you specify during inquiry definition.

The number of cycles maintained at your installation for each of the CA MICS files in Figure 2-81 determines the time frame for which you may produce ASM Analysis reports and plots.

2.4 I/O Configuration Analysis

The I/O Configuration Analysis routines allow you to analyze the following aspects of your I/O configuration:

- o Analysis of shared DASD effects for multiple systems.
- o Analysis of the components of I/O device service time to identify problems in device response time.
- o Analysis of DASD skew to identify DASD strings that are adversely impacted by high activity devices.
- o Identification of 3350/3380 volumes for back-to-back mapping on dual density devices.
- o Analysis of tape drive allocation patterns to identify problems in tape drive availability and model possible solutions.
- o Analysis of FICON channel service times and performance metrics to ensure proper channel usage.

The following sections describe the operation of the I/O Configuration Analysis inquiries and the reports they produce.

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Shared DASD Analysis Inquiry
- 4 - DASD I/O Component Analysis Inquiry
- 5 - DASD Skew Analysis Inquiry
- 6 - DASD Multiple Volume Analysis Inquiry
- 7 - Tape Drive Simulation Analysis Inquiry
- 8 - FICON Channel Performance Analysis Inquiry

2.4.1 Functional Description

The I/O Configuration Analysis inquiries produce six primary reports. Each of these reports allows you to analyze a different facet of the performance of your I/O configuration. The reports are:

- o Shared DASD Analysis Report - Details the utilization and performance impacts of DASD devices that are shared by multiple systems. A sample Shared DASD Analysis Report is shown in Figure 2-82.
- o I/O Response Time Component Analysis Report - Breaks device response time into internal queuing, queue for control unit, queue for device, and start I/O times. A sample I/O Response Time Component Analysis Report is shown in Figure 2-83.
- o DASD Skew Analysis Report - Analyzes DASD device strings to identify strings that have heavily utilized devices that adversely influence the performance of the remainder of the devices on the string. A sample DASD Skew Analysis Report is shown in Figure 2-84.
- o DASD Multiple Volume Analysis Report - Provides suggested volume groupings for DASD volumes to higher density devices. An analysis of existing DASD volumes can also be produced. A sample DASD Multiple Volume Analysis report is shown in Figure 2-85.
- o Tape Drive Simulation Analysis Report - Analyzes the utilization of the tape drive subsystem. Modeling the effects of changes in the tape drive subsystem's workload can also be done. A sample Tape Drive Simulation Analysis report is shown in Figure 2-86.
- o Channel Resource Usage Summary Report - Analyzes the performance of each FICON channel across a user provided time interval to quantify channel performance and to identify potential sources of performance problems.

Additionally, two other reports, the Processor Complex & Channel Configuration and the Channel Resource Usage Detail reports, and the Channel Service Time graphs can be optionally produced as well. Each of these reports will be discussed in more detail in a later section.

The operation and interpretation of each of these reports are discussed in detail in the remainder of this chapter. The

algorithms used to produce these reports are discussed in the technique tutorials in Sections 2.4.3.2, 2.4.4.2, 2.4.5.2, 2.4.6.2, 2.4.7.2 and 2.4.8.2.

CA MICS Performance Manager																		
SHARED DASD ANALYSIS REPORT																		
INTERVAL START: 03JANyy:23:59																		
INTERVAL END : 04JANyy:15:59																		
DURATION : 16:00																		
DEV		DEVICE	TOTAL	%	---AVERAGE TIME (MS)---				DEV	SIO	%	%	AV RESP		AV SERV		AV QUE	
ADR	VOLSER	TYPE	SIO	UTIL	RESP.	SERVICE	QUEUE	SYSID	ADR	COUNT	SIO	UTIL	TIME MS	R/R	TIME MS	R/S	TIME MS	R/Q
16C	PAG16C	3380	111E3	12.5	33.3	32.5	0.8	SYSA	16C	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	16C	111E3	100	12.5	33.3	1.0	32.5	1.0	0.8	1.0
								SYSC	16C	22	0	0.0	45.5	1.4	45.5	1.4	0.0	0.0
16D	SMF302	3380	95251	7.2	21.2	21.0	0.2	SYSA	16D	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	16D	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSC	16D	95251	100	7.2	21.2	1.0	21.0	1.0	0.2	1.0
161	PAG161	3380	83489	6.4	23.4	22.0	1.4	SYSA	161	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	161	83467	100	6.4	23.4	1.0	22.0	1.0	1.4	1.0
								SYSC	161	22	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
168	PAG168	3380	108E3	12.6	34.3	33.4	0.9	SYSA	168	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	168	108E3	100	12.6	34.3	1.0	33.4	1.0	0.9	1.0
								SYSC	168	22	0	0.0	45.5	1.3	0.0	0.0	45.5	49
169	SMF301	3380	37098	2.8	21.6	21.1	0.6	SYSA	169	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	169	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSC	169	37098	100	2.8	21.6	1.0	21.1	1.0	0.6	1.0
17D	PAG17D	3380	47671	7.0	57.0	42.6	14.5	SYSA	17D	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	17D	47649	100	7.0	57.1	1.0	42.6	1.0	14.5	1.0
								SYSC	17D	22	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
17F	M33803	3380	82955	4.4	25.5	21.3	4.3	SYSA	17F	44738	54	1.8	29.6	1.2	23.0	1.1	6.7	1.6
								SYSB	17F	20426	25	1.3	21.1	0.8	18.7	0.9	2.4	0.6
								SYSC	17F	17791	21	1.3	20.2	0.8	19.9	0.9	0.3	0.1
171	PAG171	3380	47988	6.6	52.0	39.5	12.6	SYSA	171	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	171	47966	100	6.6	52.1	1.0	39.5	1.0	12.6	1.0
								SYSC	171	22	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
173	PAG173	3380	47793	6.4	50.2	38.3	11.9	SYSA	173	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	173	47771	100	6.4	50.3	1.0	38.3	1.0	11.9	1.0
								SYSC	173	22	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
174	M33801	3380	38972	2.0	21.0	19.9	1.1	SYSA	174	20741	53	0.7	20.4	1.0	19.0	1.0	1.4	1.3
								SYSB	174	7622	20	0.5	18.8	0.9	18.1	0.9	0.7	0.6
								SYSC	174	10609	27	0.9	23.8	1.1	23.0	1.2	0.8	0.8

Figure 2-82. Sample MVS/370 Version of Shared DASD Analysis Report

Figure 2-83. Sample I/O Response Time Component Analysis Report

CA MICS Performance Manager																	
13:08 WEDNESDAY, DECEMBER 19, yyyy																	
DASD SKEW ANALYSIS - DAYS LEVEL																	
INTERVAL START: 15NOVyy:09:52																	
INTERVAL END : 15NOVyy:10:52																	
DURATION : 1:00:05																	
SYSID : SYSA																	
CU /	AV	LCU	TOTAL	TOTAL	TYPE									CV /			
PATH	LCU	SERV	MS	SSCH	RTE	%	HOS	LINE	DEV 1	DEV 2	DEV 3	DEV 4	DEV 5	DEV 6	DEV 7	DEV 8	SKEW
1A0-0 0005	10.26		33.2	34.08	DEV NO.	1A0	1A1	1A2	1A3	1A4	1A5	1A6	1A7	--SEE--	* DLS *		
					SSCH RT	0.3	0.0	0.0	0.0	0.0	0.0	0.0	0.3	-BELOW-	ANY TWO		
					AV SERV	30.63	0.00	37.06	0.00	0.00	0.00	0.00	0.00	21.60	DEVICES		
					% PATH	0.12	0.00	0.00	0.00	0.02	0.00	0.00	0.21		CAN BE		
					VOLSER	WORK21	XTG001	BENSC2	FYQ001	WORK30	WORK31	PIAN01	WMN017		ACTIVE		
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A				
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J				
1A0-0 0005	-----SEE ABOVE-----				DEV NO.	1A8	1A9	1AA	1AB	1AC	1AD	1AE	1AF	2.14	* DLS *		
					SSCH RT	0.0	0.0	0.0	31.5	0.0	0.1	0.0	1.1	1.36	ANY TWO		
					AV SERV	0.00	0.00	0.00	22.16	0.00	10.21	0.00	22.48		DEVICES		
					% PATH	0.00	0.00	0.00	32.96	0.00	0.02	0.00	0.76		CAN BE		
					VOLSER	PIAN05	FAT1B9	CIS001	FAT1BB	FAT1BC	FAT1BD	FAT1BE	TECSYS		ACTIVE		
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A				
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J				
1B0-0 0005	9.30		15.7	14.60	DEV NO.	1B0	1B1	1B2	1B3	1B4	1B5	1B6	1B7	--SEE--	* DLS *		
					SSCH RT	.15	0.4	0.0	0.4	0.1	0.0	3.3	0.0	-BELOW-	ANY TWO		
					AV SERV	10.78	22.82	0.00	23.73	15.84	18.48	16.26	16.53		DEVICES		
					% PATH	0.32	0.53	0.00	0.21	0.02	0.00	0.00	0.00		CAN BE		
					VOLSER	WORK22	WORK26	WORK23	CH0100	PIAN02	PIAN04	BES003	BES001		ACTIVE		
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A				
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J				
1B0-0 0005	-----SEE ABOVE-----				DEV NO.	1B8	1B9	1BA	1BB	1BC	1BD	1BE	1BF	2.25	* DLS *		
					SSCH RT	10.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.13	ANY TWO		
					AV SERV	24.30	0.00	48.81	0.00	0.00	0.00	0.00	0.00		DEVICES		
					% PATH	12.71	0.00	0.02	0.00	0.00	0.00	0.00	0.00		CAN BE		
					VOLSER	WMN003	WMN004	IPN001	WMN006	PIAN03	IPN002	CS0101	PIAN06		ACTIVE		
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A				
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J				
1C0-0 0006	3.51		30.2	10.56	DEV NO.	1C0	1C1	1C2	1C3	1C4	1C5	1C6	1C7	--SEE--	* DLS *		
					SSCH RT	0.0	0.0	0.1	0.1	0.5	8.5	2.1	0.6	-BELOW-	ANY TWO		
					AV SERV	15.81	0.00	38.00	27.63	15.84	18.48	16.26	16.53		DEVICES		
					% PATH	0.00	0.00	0.11	0.08	0.10	3.30	0.41	0.15		CAN BE		
					VOLSER	SPL007	SPL008	WMN025	WMN026	WJEP03	WJEP04	FEJ000	DLIB23		ACTIVE		
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A				
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J				
1C0-0 0006	-----SEE ABOVE-----				DEV NO.	1C8	1C9	1CA	1CB	1CC	1CD	1CE	1CF	1.24	* DLS *		
					SSCH RT	0.0	0.0	3.0	0.0	12.0	1.3	0.0	2.1	1.69	ANY TWO		
					AV SERV	20.60	28.42	31.04	35.61	19.62	17.51	0.00	21.04		DEVICES		
					% PATH	0.00	0.00	1.90	0.00	3.29	0.32	0.00	0.90		CAN BE		
					VOLSER	TEC003	WJEP18	IPN004	MICS02	FAT25C	FAT25D	WMN012	WMN013		ACTIVE		
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A				
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J				

Figure 2-84. Sample DASD Skew Analysis Report

CA MICS Performance Manager					
DASD MULTIPLE VOLUME ANALYSIS					
INTERVAL START: 30SEPyy:07:30					
INTERVAL END : 07SEPyy:17:59					
DURATION : 26:00					
SYSIDS : A370 B370 XA01 XA02					
VOLSER1	DEVICE TYPE1	VOLSER2	DEVICE TYPE2	AVERAGE PRODUCT	MAXIMUM SUM
MIS380	3380	PAG104	3380	0.20	1.69
LIB503	3380	PAG103	3380	2.18	5.49
PROD01	3350	TSSUCC	3350	8.28	14.01
HSMCTL	3380	SP30L3	3380	35.84	69.13

Figure 2-85. Sample DASD Multiple Volume Analysis Report

CA MICS Performance Manager FICON Channel Performance Analysis Channel Resource Usage Summary Report, by SYSID For Friday, September 3, yyyy																								
Ending Date: 20JULyyyy					Duration: 24:17:25																			
CEC Serial: 0851 LCSS: 0																								
CHPID: 42 Description: FICON Switched					DCM in Use: Y CHPID Group: 2 Shared: Y																			
Channel Usage: DASD																								
BUS Usage Count: 739.8M					BUS Util. Pct : 13.09					Total Work Units : 2.5G					CHPID Pct Busy : 22.41					CHPID I/Os p/sec : 362				
Bytes Read p/Sec: 5.1M					Read Path % Busy : 4.27					Bytes Writ p/Sec: 3.3M					Write Path % Busy: 2.81									
Avg Burst Size: 23.4K					Avg Svc Time (ms): 2.40					Est. Open Exchngs: 0.80					Estimated Credits: 10									
LPAR Usage Stats:																								
SYSID	Work Units	CHPID % Busy	Total Bytes Read	Bytes Read p/sec	Read Path % Busy	Total Bytes Written	Bytes Written p/sec	Write Path % Busy	Total I/O Count	I/Os p/sec	Avg Burst Size	Average Svc Time (ms)	Avg Open Xchngs	Estimd. Credits										
CA11	80M	0.71	8G	98.3K	0.08	11G	125.9K	0.10	1M	19	11.2K	2.39	0.04	0										
CA31	566M	4.97	136G	1.5M	1.30	132G	1.5M	1.26	5M	68	44.8K	1.34	0.08	2										
CA71	20M	0.18	2G	25.5K	0.02	4G	47.5K	0.04	615K	7	10.3K	1.91	0.01	0										
XAD1	1G	13.81	231G	2.6M	2.21	98G	1.1M	0.94	20M	240	15.7K	2.66	0.59	5										
XAE1	21M	0.19	4G	49.0K	0.04	5G	60.8K	0.05	392K	4	24.5K	2.86	0.01	0										
XAT1	12M	0.11	7G	91.1K	0.08	1G	19.3K	0.02	199K	2	48.3K	3.67	0.01	0										
XA55	6M	0.06	1G	12.1K	0.01	2M	26	0.00	222K	2	4.8K	2.41	0.01	0										
XE21	39M	0.35	10G	115.2K	0.10	9G	109.8K	0.09	453K	5	43.3K	2.12	0.01	0										
XE44	40M	0.36	15G	176.6K	0.15	6G	76.0K	0.06	881K	10	25.0K	3.49	0.03	0										
XE61	64M	0.56	17G	194.7K	0.16	18G	209.5K	0.17	127K	1	277.1K	2.74	0.00	0										

Figure 2-86A. Sample Channel Resource Usage Summary Report

2.4.2 Usage Guidelines

The primary consideration for the use of the I/O Configuration Analysis inquiries is the volume of data specified for processing by each of the analytic routines. We recommend the following guidelines for using the I/O Configuration Analysis inquiries:

- o Use Shared DASD Analysis inquiry to analyze short peak intervals for a single day or over several days. Large volumes of data should not be processed, since pronounced shared DASD problems that may exist over a short period of time can be lost when averaged in with a significant number of hours in which no problems existed. For example, the inquiry is best used to analyze the impact of shared DASD during the execution of online systems or during long-running batch jobs.
- o Use the DASD I/O Component Analysis inquiry to analyze short peak intervals for a single day. Once again, the problems that this algorithm identifies are easily washed out by the summarization of multiple hours.
- o Use the DASD Skew Analysis inquiry to analyze one- or two-hour periods. DASD skew is an extremely transient problem that can be hidden by the summarization of even a few hours.
- o Provide the DASD Multiple Volume Analysis inquiry with the maximum possible volume of input data. Before mapping volumes together, you should analyze the consequences of the decision in every possible environment.
- o Use the Tape Drive Simulation Analysis inquiry to analyze short peak intervals for a single day or over several days. Large volumes of data should not be processed over long time intervals, since pronounced tape drive utilization problems that may exist over a short period of time can be lost when averaged in with a significant number of hours in which no problems exist. The inquiry is best used to analyze the impact of tape drive utilization during long-running batch jobs or during periods when multiple batch jobs may be constrained by the availability of tape drives.

Although you may wish to evaluate other quantities of data, we suggest that you follow the recommendations made in this section until you are familiar with the interpretation of the

reports.

Figure 2-87 shows the I/O Configuration Analysis Menu, which you can use to select the various I/O Configuration Analysis reports.

```
/----- Input/Output Configuration Analysis (IOA) -----\
| Option ==>                                              |
| Composing CA MICS Inquiry: IOA01 - Volume Migration Study |
|   1 - Shared DASD Analysis                               |
|       Analysis of the effects of shared DASD usage.      |
|   2 - DASD I/O Component Analysis                       |
|       Analysis of the component parts of DASD response time. |
|   3 - DASD Skew Analysis                                 |
|       Analysis of string utilization.                     |
|   4 - Multiple DASD Volume Analysis                     |
|       Analysis of DASD volume combination on activity.   |
|   5 - Tape Drive Simulation Analysis                    |
|       Simulation of tape drive utilization by program.    |
|-----/
```

Figure 2-87. I/O Configuration Analysis Menu

2.4.3 Shared DASD Analysis Inquiry

This section addresses the utilization of DASD devices that are shared by multiple systems. The following topics are discussed:

- 1 - Standard Output
- 2 - Technique Tutorial
- 3 - Case Study
- 4 - Component Operation

2.4.3.1 Standard Output

The Shared DASD Analysis Report shown in Figure 2-88 provides you with a view of the total utilization of each of the DASD devices by all of the systems in the hardware configuration.

The report details how the performance of the device on each of the sharing systems compares to the overall average performance of the device for all systems in the configuration. Modifications of the report format for MVS/XA are discussed below. The Shared DASD Analysis Report contains the following information:

INTERVAL START:	The time and date of the first HARDVA File observation included in the analysis.
INTERVAL END:	The time and date of the last HARDVA File observation included in the analysis.
DURATION:	The total number of hours summarized by report, rounded up to the nearest hour.
DEV ADR/ DEV NO.:	The address of the device on the primary system, that is, the system defined by the SYSID control parameter. The same device need not be designated by the same device address by each system in the configuration. For MVS/XA and MVS/ESA this report column is replaced by the value DEV NO, device number.
VOLSER:	The volume serial number of the volume.
DEVICE TYPE:	The device type as carried in the CA MICS DEVTYPE data element.
TOTAL SIO/ TOT SSCH:	The total number of start I/Os (SIOs) or start subchannels issued to the device by all of the systems in the configuration.
% UTIL:	The total utilization of the device by all of the systems in the configuration. Each system's utilization is calculated by taking either the maximum of the device percent busy (DVAPCACT) or device percent reserved (DVAPCRSV), that is, the maximum of the percent of time a system used a device or reserved the device so that other systems could not access it.

RESP:	The average response time (in milliseconds) for the device for all systems in the configuration. This value includes both device busy time and internal queuing for the device. Average response time is CA MICS data element DVAVRTM.
SERVICE:	The average service time (milliseconds) for the device for all systems in the configuration. This value is computed based on the utilization of the device and the number of SIOs completed. Average service time is CA MICS data element DVAVSTM.
QUEUE:	The average internal queuing delay (milliseconds) before IOS is able to start an I/O to the system. Queue time is the difference between device response and service time. Typically, queue time increases with device utilization.

For each system that shares the device, one line is produced with the following format:

SYSID:	The SYSID of the system that is to be analyzed that shares the device.
DEV ADR/ DEV NO.:	The device address or number by which the device is known on the specified SYSID. As was mentioned earlier, the same device need not be designated by the same device address by each system in the configuration. However, in this example, all of the systems use identical addresses for each of the devices.
SIO COUNT/ SSCH CNT:	The number of start I/Os or start subchannels issued to the device by the specified SYSID.
% SIO:	The percent of the total SIOs issued by the specified SYSID.
% UTIL:	The percent utilization of the device by the specified SYSID.
AV RESP TIME MS:	The average response time (milliseconds) for the device on the specified system.
R/R:	The ratio of the device response time on the system to the average device response time for the configuration. Values less than or equal to

one indicate that a system is attempting to monopolize a device. Values greater than one indicate that a SYSID is being impacted by the other systems' utilizations.

AV SERV
TIME MS: The average service time (milliseconds) for the device on the specified system.

R/S: The ratio of the device service time on the system to the average device service time for the configuration. Values less than or equal to one indicate that a system is attempting to monopolize a device. Values greater than one indicate that a SYSID is being impacted by the other systems' utilizations. For a well-balanced configuration, all systems should either have values of zero (indicating no utilization of the device) or values between 0.9 and 1.1.

AV QUE
TIME MS: The average internal queue time (milliseconds) for the device on the specified system.

R/Q: The ratio of the device internal queue time on the system to the average device internal queue time for the configuration. Values less than one indicate that the device is lightly utilized on the system. Values greater than one indicate that the internal queuing delay on the system may be substantial.

CA MICS Performance Manager																	
SHARED DASD ANALYSIS REPORT																	
INTERVAL START: 29SEPy:08:00																	
INTERVAL END : 29SEPy:16:59																	
DURATION : 9:00:00																	
DEV NO.	VOLSER	DEVICE TYPE	TOTAL SSCH	% UTIL	---AVERAGE TIME (MS)---			DEV SYSID	SSCH COUNT	% SIO	% UTIL	AV RESP TIME MS R/R		AV SERV TIME MS R/S		AV QUE TIME MS R/Q	
104	HSMCTL	3380	414K	23.3	14.7	14.5	0.2	XA01 104	399K	96	21.7	14.6	1.0	14.4	1.0	0.2	0.8
								XA02 104	15146	4	1.6	18.6	1.3	16.9	1.2	1.7	7.4
110	PAG303	3380	127K	13.1	29.6	29.6	0.0	XA01 110	127K	100	13.1	29.6	1.0	29.6	1.0	0.0	0.0
								XA02 110	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
114	PAG302	3380	272K	24.3	25.9	25.9	0.0	XA01 114	272K	100	24.3	25.9	1.0	25.9	1.0	0.0	0.0
								XA02 114	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
120	WORK13	3380	101K	9.2	31.9	27.8	4.1	XA01 120	101K	100	9.2	31.9	1.0	27.8	1.0	4.1	1.0
								XA02 120	56	0	0.0	16.4	0.5	16.4	0.6	0.0	0.0
123	TS0805	3380	97510	4.9	11.9	11.3	0.6	XA01 123	79266	81	3.5	12.2	1.0	11.4	1.0	0.8	1.2
								XA02 123	18244	19	1.4	10.7	0.9	10.7	1.0	0.0	0.0
132	WORK22	3380	406K	27.5	18.8	17.2	1.6	XA01 132	352K	87	20.4	17.9	1.0	16.7	1.0	1.2	0.8
								XA02 132	53388	13	7.1	24.1	1.3	20.3	1.2	3.9	2.4
150	WORK15	3380	83256	6.3	26.2	22.9	3.3	XA01 150	83256	100	6.3	26.2	1.0	22.9	1.0	3.3	1.0
								XA02 150	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
160	WORK12	3380	107K	7.9	25.1	23.2	1.9	XA01 160	107K	100	7.9	25.1	1.0	23.2	1.0	1.9	1.0
								XA02 160	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
170	WORK14	3380	94524	7.0	24.3	22.8	1.4	XA01 170	94524	100	7.0	24.3	1.0	22.8	1.0	1.4	1.0
								XA02 170	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
225	LIBS03	3380	79290	6.1	20.1	18.6	1.5	XA01 225	58820	74	3.6	20.3	1.0	18.7	1.0	1.6	1.1
								XA02 225	20470	26	2.4	19.5	1.0	18.5	1.0	1.0	0.7
226	MCAT01	3380	77024	3.2	12.1	11.4	0.6	XA01 226	77024	100	3.2	12.1	1.0	11.4	1.0	0.6	1.0
								XA02 226	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
260	GSL802	3380	354K	19.4	18.6	15.9	2.7	XA01 260	354K	100	19.4	18.6	1.0	15.9	1.0	2.7	1.0
								XA02 260	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
305	BCSACT	3380	141K	9.5	34.4	19.5	14.9	XA01 305	141K	100	9.5	34.4	1.0	19.5	1.0	14.9	1.0
								XA02 305	324	0	0.0	19.3	0.6	19.3	1.0	0.0	0.0
310	WORK11	3380	150K	12.9	31.0	26.3	4.7	XA01 310	150K	100	12.9	31.0	1.0	26.3	1.0	4.7	1.0
								XA02 310	10	0	0.0	34.5	1.1	34.5	1.3	0.0	0.0
350	SP3CKP	3380	72184	29.4	30.0	30.0	0.0	XA01 350	55552	77	21.8	26.0	0.9	26.0	0.9	0.0	0.0
								XA02 350	16632	23	7.7	43.3	1.4	43.3	1.4	0.0	0.0
56E	M33809	3380	236K	14.0	19.5	17.1	2.4	XA01 56E	236K	100	14.0	19.5	1.0	17.1	1.0	2.4	1.0
								XA02 56E	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Figure 2-88. MVS/XA Version of Shared DASD Analysis Report

2.4.3.2 Technique Tutorial

DASD hardware allows multiple systems to be connected to a single volume. This feature is highly desirable if volumes containing critical applications must be reached from another system in the event of a processor failure, or to allow a single copy of critical data to be shared by multiple systems.

However, shared DASD often introduces significant performance problems by allowing workloads on one system to adversely affect activities on other systems. Since RMF reports on the I/O activities of individual systems rather than the I/O activity of an installation as a whole, it is necessary to summarize the RMF data from multiple systems to evaluate the performance of a shared device.

Unfortunately, RMF does not contain data elements that indicate whether or not a device is shared. The algorithm implemented by the Shared DASD Analysis inquiry assumes that disk volume serial numbers are unique within the installation. Using this assumption, all of the records (for the same time interval) for each volume serial number are mapped together for the analysis.

The shared DASD analysis algorithm is based on data contained in the CA MICS HARDVA File. The following values are computed for each shared volume:

TOTAL SIOs: The sum of the SIO count values, CA MICS data element DVASIO, for each of the SYSIDs in the configuration.

AV RESPONSE TIME: The weighted average of the response times, data element DVAAVRTM, for all of the SYSIDs that share the device. Response time is the sum of average device service time and the internal queuing for the device. Internal queuing is defined as the time the I/O waited in the unit control block UCB and logical channel queues. This value is calculated by weighting the response time measured for each of the systems by the number of I/Os issued to the device by the system.

AV SERVICE TIME: The weighted average of the service times, data element DVAAVSTM, for all of the SYSIDs that share the device. The average service time is computed by dividing the device utilization

by the start I/O rate. The value is calculated by weighting the service time measured for each of the systems by the number of I/Os issued to the device by the system.

AV QUEUE TIME: The weighted average of the internal queue times for the device (the difference of CA MICS data elements DVAVRTM and DVAVSTM) for all of the SYSIDs that share the device. The value is calculated using the weighted average service and response times.

% UTILIZATION: The sum of the percentages of the times that each of the systems utilizes the device. Prior to MVS/XA, utilization was defined as the maximum of either the percent of the time that a system used a device (i.e., device active, DVAPCACT) or the percent of the time that the system prevented other systems from using the device (i.e., device reserved, DVAPCRSV). This definition was required, since RMF did not indicate the percent of time that a reserve condition existed when no I/Os were outstanding.

Starting with MVS/XA, a new variable called DVAUTL was introduced that indicates the number of samples during which the device was reserved and no I/O was outstanding. Hence, utilization in an MVS/XA environment is defined as DVAPCACT plus the percent of time the DVAUTL condition existed.

The effect of the DASD sharing on each of the devices is determined by calculating the following ratios:

- o The ratio of each system's average response time to the weighted average response time for the configuration.
- o The ratio of each system's average service time to the weighted average service time for the configuration.
- o The ratio of each system's average internal queue time to the weighted average internal queue time for the configuration.

Each of these ratios is discussed in the following paragraphs.

The ratio of the device response time on the system to the average device response time for the configuration indicates whether I/Os issued by the system are receiving better or worse service than the average for the configuration. Since this value includes internal queuing times, it is an indicator, but not the best index, for evaluating shared DASD contention.

The ratio of the device service time on the system to the average device service time for the configuration is the best indicator of shared DASD contention. It indicates whether the system is impacting or being impacted by the I/Os issued by the other systems. Values less than or equal to one indicate that a system is attempting to monopolize a device. Values greater than one indicate that a SYSID is being impacted by the other system. For a well-balanced configuration, all systems should either have values of zero (indicating no utilization of the device) or values between 0.9 and 1.1.

The ratio of the device internal queue time on the system to the average device internal queue time for the configuration provides a relative indication of internal queuing on each of the systems. Values less than one indicate that the device is lightly utilized on the system, while values greater than one indicate that the internal queuing delay on the system may be substantial.

2.4.3.3 Case Study

This case study details the application of the I/O Configuration Analysis inquiries to produce the Shared DASD Analysis Report. In particular, the following topics are addressed:

- 1 - Problem Description
- 2 - Inquiry Panel
- 3 - Output Report

2.4.3.3.1 Problem Description

Shared DASD allows multiple systems to be connected to a single volume. This feature is highly desirable for some applications' volumes, since critical applications may be moved to another system in the event of a processor failure. Also, shared DASD allows a single copy of critical data to be shared by multiple systems.

However, shared DASD can introduce significant performance problems by allowing workloads on one system to adversely affect activities on other systems. Since RMF reports on the I/O activities of individual systems, rather than the I/O activity of an installation as a whole, it is necessary to summarize all of the HARDVA data for every system that shares a volume with the specified system to obtain an accurate view of each device.

For this case study, we used the Shared DASD Analysis Report to analyze a configuration consisting of three IBM 3081s. The objectives of the study were to determine the degree to which the shared volumes were utilized by the processors and the effect of the multisystem contention on the I/O response, service, and queue times (CA MICS data elements DVAAVRTM, DVAAVSTM, and DVAAVRTM-DVAAVSTM) of each of the systems.

The inquiry panels used to conduct the study and the output reports are discussed in the following two sections.

2.4.3.3.2 Inquiry Panel

We used I/O Configuration Analysis routines to study the impact of the shared DASD effects on system SYSA, one of three IBM 3081s that shared many of the DASD volumes in the configuration. (NOTE: Section 2.4.3.4.1 presents a detailed description of the inquiry fields on the IOA Shared DASD Analysis Report inquiry panel.)

We selected the peak morning and afternoon periods (hours 9, 10, 14, and 15) for the study over a four-day period. We denoted the CA MICS database unit that contained the installation's hardware data by database identifier B. The inquiry fields for the study are shown in Figure 2-89. The generated report is shown in Figure 2-90.

```

/----- IOA Shared DASD Analysis Report -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: IOA08 - Shared DASD Case Study |
| Time-span ==> DETAIL (DETAIL/DAYS)                Cycle(s) ==> 01 - 04 |
| Data base ==> B (PRIMARY) _                      _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
| Shared System (All sharing systems will be reported) ==> SYSA |
| Report by time interval ==> N (Y/N) Minimum I/O rate ==> 1 (.001 - 999) |
| Usage qualification processing: |
| CA MICS Element to test==> _____ (Any numeric HARDVA element) |
| Operator to employ ==> __ (LE,LT,GT,GE) |
| Numeric value to test ==> _____ (.00000000 - 99999999) |
| Data Selection: |
| Date ==> _____ - _____ - _____ ==> N (Y/N) |
| Hour ==> 09 - 10 14 - 15 __ - __ ==> N (Y/N) |
| Zone ==> _ _ _ _ _ ==> N (Y/N) |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-89. Case Study - IOA Shared DASD Analysis Report Inquiry Panel

2.4.3.3.3 Output Report

Figure 2-90 shows the Shared DASD Analysis Report resulting from the information entered on the IOA Shared DASD Analysis Report inquiry panel. There are a number of values in the report that you should carefully examine.

- o The total percent utilization for all systems in the configuration, that is, the left-most column labeled % UTIL. If this value is in excess of 30-35 percent, then significant problems can be expected both in internal queuing and average response time. The average response time will increase significantly, due to significant internal queuing for the device.
- o High average internal queue time (the center column labeled QUEUE). This is the symptom of multiple address spaces in a single system competing for a single device. You should examine the contents of the volume to see if the volume could be split into two separate volumes on different paths.
- o Utilization of a shared device by a single system. This is shown most graphically by the %SIO column. If a device is consistently utilized by only a single system, you should raise the question as to why the volume is physically shareable (or at least online) to the other system. This problem can be corrected by normally varying it offline or by making it physically non-shareable.
- o High values (that is, greater than 1.1) of the R/R and R/S columns. These values indicate systems whose performance is being significantly impacted by the other systems that share the device.

By reviewing these values, you can identify devices where the effects of shared DASD are having a significant impact. Generally, you should ask the question, "why should this device be shared?" rather than, "why not share the device?"

CA MICS Performance Manager																			
SHARED DASD ANALYSIS REPORT																			
INTERVAL START: 27DECyy:15:44																			
INTERVAL END : 04JAN91:15:59																			
DURATION : 16:00																			
DEV		DEVICE	TOTAL	%	---AVERAGE TIME (MS)---			SYSID	DEV	SIO	%	%	AV RESP		AV SERV		AV QUE		
ADR	VOLSER	TYPE	SIO	UTIL	RESP.	SERVICE	QUEUE		ADR	COUNT	SIO	UTIL	TIME MS	R/R	TIME MS	R/S	TIME MS	R/Q	
16C	PAG16C	3380	111E3	12.5	33.3	32.5	0.8	SYSA	16C	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	16C	111E3	100	12.5	33.3	1.0	32.5	1.0	0.8	1.0	
								SYSC	16C	22	0	0.0	45.5	1.4	45.5	1.4	0.0	0.0	
16D	SMF302	3380	95251	7.2	21.2	21.0	0.2	SYSA	16D	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	16D	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSC	16D	95251	100	7.2	21.2	1.0	21.0	1.0	0.2	1.0	
161	PAG161	3380	83489	6.4	23.4	22.0	1.4	SYSA	161	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	161	83467	100	6.4	23.4	1.0	22.0	1.0	1.4	1.0	
								SYSC	161	22	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
168	PAG168	3380	108E3	12.6	34.3	33.4	0.9	SYSA	168	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	168	108E3	100	12.6	34.3	1.0	33.4	1.0	0.9	1.0	
								SYSC	168	22	0	0.0	45.5	1.3	0.0	0.0	45.5	49	
169	SMF301	3380	37098	2.8	21.6	21.1	0.6	SYSA	169	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	169	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSC	169	37098	100	2.8	21.6	1.0	21.1	1.0	0.6	1.0	
17D	PAG17D	3380	47671	7.0	57.0	42.6	14.5	SYSA	17D	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	17D	47649	100	7.0	57.1	1.0	42.6	1.0	14.5	1.0	
								SYSC	17D	22	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
17F	M33803	3380	82955	4.4	25.5	21.3	4.3	SYSA	17F	44738	54	1.8	29.6	1.2	23.0	1.1	6.7	1.6	
								SYSB	17F	20426	25	1.3	21.1	0.8	18.7	0.9	2.4	0.6	
								SYSC	17F	17791	21	1.3	20.2	0.8	19.9	0.9	0.3	0.1	
171	PAG171	3380	47988	6.6	52.0	39.5	12.6	SYSA	171	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	171	47966	100	6.6	52.1	1.0	39.5	1.0	12.6	1.0	
								SYSC	171	22	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
173	PAG173	3380	47793	6.4	50.2	38.3	11.9	SYSA	173	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
								SYSB	173	47771	100	6.4	50.3	1.0	38.3	1.0	11.9	1.0	
								SYSC	173	22	0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
174	M33801	3380	38972	2.0	21.0	19.9	1.1	SYSA	174	20741	53	0.7	20.4	1.0	19.0	1.0	1.4	1.3	
								SYSB	174	7622	20	0.5	18.8	0.9	18.1	0.9	0.7	0.6	
								SYSC	174	10609	27	0.9	23.8	1.1	23.0	1.2	0.8	0.8	

Figure 2-90. Case Study - Shared DASD Analysis Report

2.4.3.4 Component Operation

The inquiry fields let you define the CA MICS database cycles to be used as input, the time periods to be analyzed, and the reports to be produced. Figures 2-91 through 2-95 show the control panels used to produce the IOA Shared DASD Analysis Report.

The following sections describe each control panel:

- 1 - Shared DASD Analysis Report Inquiry Panel
- 2 - Extended Options Menu Panel
- 3 - Global Data Selection Panel
- 4 - Alternate Input File Panel
- 5 - Save Selected Input Observations Panel

2.4.3.4.1 Shared DASD Analysis Report Inquiry Panel

Figure 2-91 shows the IOA Shared DASD Analysis Report inquiry panel that you use to prepare Shared DASD Analysis Reports.

The following is a discussion of the fields on this panel.

Timespan:	The specification of whether the data is to be taken from the DETAIL or DAYS CA MICS timespan.
Cycle(s):	One or more cycles of the CA MICS files that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a cycle of the CA MICS file usually contains a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided by the DATE and HOUR control parameters.
Database:	The single-character CA MICS identifiers for the databases you want to access.
Allow Runtime Override:	The indication of whether the database identifier and/or the Cycles specification can be overridden at inquiry execution time.
Shared System:	The SYSID of the data to be reported.
Report by time interval:	The specification of whether a report should be generated for each interval specified or summarized into a single report.
Minimum I/O rate:	A filter for the exclusion of low I/O activity volumes. It specifies the minimum per second I/O rate for volume inclusion in the analysis. This value may range between 0.001 and 999. This allows you to exclude devices that were either unused during the interval or received very little use. If you use this control parameter, the I/O filter applies to all I/O component and shared-DASD reports.
CA MICS Element to test:	In addition to the Minimum I/O rate, the Shared DASD Analysis Report inquiry can test any other numeric data element in the HARDVA File for a specific condition. This can be used to further subset the data to be analyzed.

Operator The boolean operator to employ - LE, LT, GT,
to employ: or GE.

Numeric
value to The value against which the data element is
test: tested.

Date: The day(s) to be used for analysis. This field
is normally used with the IOA Shared DASD
Analysis Report inquiry to drop the extra day
from a cycle that is due to the timing of the
SMF dump.

Hour: An hour or range of hours for which data is
reported.

Zone: The specification that zone selection is to be
applied. If zone 1 definition is 8:00 a.m. to
6:00 p.m., daily reports covering that time
period can be specified by using the ZONE
field set to 1.

Specify
extended The specification that allows you to supply
options: additional report options. These options
include special SAS data selection statements,
non-standard input, and output files. These
options are discussed in the sections that
follow.

The operation of the time selection control parameters depends on definitions made during the installation of CA MICS. If you are not sure about the zone definitions for your installation or the CA MICS database update procedures, consult your CA MICS System Administrator.

```

/----- IOA Shared DASD Analysis Report -----\
| Command ==>                                     |
|                                                    |
| Composing CA MICS Inquiry: IOA10 - Shared DASD SYS1 |
|                                                    |
| Time-span ==> DETAIL (DETAIL/DAYS)                Cycle(s) ==> 01 - 02 |
| Data base ==> A (PRIMARY) _                      _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
|                                                    |
| Shared System (All sharing systems will be reported) ==> SYS1 |
| Report by time interval ==> N (Y/N) Minimum I/O rate ==> 1 (.001 - 999) |

```

```
|
| Usage qualification processing:
| CA MICS Element to test==> _____ (Any numeric HARDVA element)
| Operator to employ      ==>  __ (LE,LT,GT,GE)
| Numeric value to test  ==>  _____ (.00000000 - 99999999)
|
| Data Selection:                                     Allow run-time override
| Date      ==>  _____ - _____ - _____ ==> N (Y/N)
| Hour      ==>  __ - __ - __ - __ - __ - __ ==> N (Y/N)
| Zone      ==>  - - - - - - - - - - ==> N (Y/N)
|
| Specify extended options ==> N (Y/N)
|
|-----/
```

Figure 2-91. Case Study - IOA Shared DASD Analysis Report Inquiry Panel

2.4.3.4.2 Extended Options Menu Panel

If you enter Y (Yes) in the Specify extended options field, the Extended Options Menu, shown in Figure 2-92, is displayed.

[illegible]

Figure 2-92. Case Study - IOA Shared DASD Analysis Report Extended Options Menu Panel

This Extended Options Menu panel allows you to specify that:

- o Global data selection statements are to be defined.
- o Non-CA MICS input files are to be defined.
- o Non-CA MICS output files are to be defined.

The default for all three options is N (No). The R (Reset) option allows you to return the option to its original contents. For detailed information about the extended options, press the help key (usually PF1) when you are at the Extended Options Menu.

2.4.3.4.3 Global Data Selection Panel

If you enter Y (yes) in the Specify global data selection exit field, the Global Data Selection Panel, which is shown in Figure 2-93, is displayed. This panel allows you to enter SAS statements that will perform special data selection functions in addition to the normal DATE/HOUR/ZONE option on the main inquiry panel.

```

/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step: IOA Shared DASD Analysis                |
| Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy  |
| Cmd SAS Statements for: Specifying global data selection |
| - ----- |
| - %MACRO SELECT; |
| - IF SYSID='SYSA' THEN DO; |
| -   IF VOLSER='SYSRES' THEN VOLSER='SYSAAA'; |
| - END; |
| - IF SYSID='SYSB' THEN DO; |
| -   IF VOLSER='SYSRES' THEN VOLSER='SYSBBB'; |
| - END; |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - %MEND SELECT; |
| ***** BOTTOM OF DATA ***** |
\-----/

```

Figure 2-93. Case Study - IOA Shared DASD Analysis Report Global Data Selection Panel

To illustrate the use of this option, consider an installation that is comprised of two processors, SYSA and SYSB. Although most of the volume serial numbers in the installation are unique, the system resident volumes of each of the systems are named SYSRES. To prevent the I/O Shared DASD Analysis inquiry from assuming that the two unique SYSRES packs are only one shared volume, you would code this Global Data Solution Exit shown in Figure 2-93.

Unless you specify the SELECT MACRO, the assumptions made by the I/O Analysis routines may produce invalid results.

2.4.3.4.4 Alternate Input File Panel

If you enter Y (Yes) in the Specify alternate input file field in the Extended Options Menu, the Alternate Input File panel, shown in Figure 2-94, is displayed.

```
/----- Alternate Input File -----  
| Command ==>  
|  
| Inquiry Step: IOA Shared DASD Analysis  
|  
| Specified Time-span: DETAIL  
|  
| Equivalent CA MICS file(s): HARDVA  
| Alternate SAS file(s)    ==> _____  
|  
| SAS data set allocation instructions:  
|   Data set name          ==> _____  
|   Disposition            ==> SHR (SHR/OLD)  
|  
|  
|  
|  
|  
|  
|  
|  
|  
|
```

Figure 2-94. Case Study -IOA Shared DASD Analysis Report Alternate Input File Panel

This panel allows you to specify a SAS file that is not a part of the standard CA MICS online database as input to the IOA Analysis routines. You must specify one or more indicated SAS file names and an MVS data set name. If you do not enclose the data set name in single quotes, your TSO prefix (or user ID) is added as the first portion of the data set name.

2.4.4 DASD I/O Component Analysis Inquiry

The DASD I/O Component Analysis Report breaks down DASD response time into its component parts. The report also displays usage and activity measures to assist you in determining the impact of response time on system performance.

This section discusses the following topics:

- 1 - Standard Output
- 2 - Technique Tutorial
- 3 - Case Study
- 4 - Component Operation

2.4.4.1 Standard Output

I/O Component Analysis produces a report that allows you to apply the methodology discussed in Section 2.4.4.2, Technique Tutorial, to analyze the individual device response time characteristics of your system.

MVS/XA AND MVS/ESA I/O COMPONENT ANALYSIS REPORT

A sample of the I/O Component Analysis report for MVS/XA and MVS/ESA is shown in Figure 2-96. All time measurements for the components of I/O in the report are in milliseconds. The following information is shown in the figure:

INTERVAL START:	The time and date of the first HARDVA File observation included in the analysis.
INTERVAL END:	The time and date of the last HARDVA File observation included in the analysis.
DURATION:	The total number of hours summarized by report, rounded up to the nearest hour.
SYSID:	The system identifier.
DEV NO.:	The device number, CA MICS data element DEVADDR.
VOLSER:	The volume serial number of the volume mounted on the device.
DEVICE TYPE:	The device type as carried by the CA MICS DEVTYPE data element.
SSCH RATE:	The number of start subchannels commands per second (SSCH) issued to the device by the system.
% BUSY:	The percent device active (DVAPCACT) as measured by RMF.
RESP:	The estimated average response time. This value is calculated as the sum of average queueing delay, pending time, disconnect time, connect time as described below.
QUEUE:	The estimated average internal queueing delay before IOS is able to start an I/O to the device. This value is derived from $(DVAQUE/SAMPLES) / (DVASSC/DURATION) * 1000$.

This can be viewed as the measured probability of the device being queued multiplied by the average length of time between I/O events.

- PEND: The measured average pending time for the device computed from data elements (DVAPEN/DVAMEC) * 1000. This value is measured by the hardware in an MVS/XA or MVS/ESA environment.
- DISC: The measured average disconnect time for the device computed from data elements (DVADIS/DVAMEC) * 1000. This value includes seek, latency, and RPS miss delays. It is measured by the hardware in an MVS/XA or MVS/ESA environment.
- CONN: The measured average device I/O connect time computed from data elements (DVACNN/DVAMEC) * 1000. This value includes CCW fetch and data transfer. It is measured by the hardware in an MVS/XA or MVS/ESA environment.
- ESCON: The measured average time that a device was delayed due to ES connection (ESCON) director port busy, computed from data elements (DVADPBTM/DVAMEC) * 1000. This value is measured by the hardware in an MVS/ESA environment with Enterprise System connection directors present.

To assist in the interpretation of the components of I/O, the values of the I/O component times are plotted over a range of 55 milliseconds. This plot is presented on the right-hand side of the report. The characters used to denote the I/O components are:

- Q - The internal queuing delay time for the device.
- P - The pending delay time for the device.
- D - The device disconnect time. This value includes seek, latency, and RPS miss delays.
- C - Connect time. This value includes CCW fetch and data transfer.

Each character represents one millisecond. The calculated values are rounded up by adding 0.49 milliseconds to the calculated values. Values exceeding 55 milliseconds are

MVS/370 I/O COMPONENT ANALYSIS REPORT

A sample of the I/O Component Analysis Report for an MVS/370 system is provided in Figure 2-97. All time measurements for the components of I/O in the report are in milliseconds. The information in the MVS/370 report that is measured or calculated differently from the information on the MVS/XA and MVS/ESA report is listed below:

DEV ADR:	The device address in data element DEVADDR.
SIO RATE:	The number of start I/Os per second (SIOs) issued to the device by the system. Note that the program excludes devices that have an SIO rate of less than one per second, because the equations that approximate the I/O time components are poorly behaved for lightly loaded devices.
RESP:	The estimated average response time. This value is calculated from the device utilization, average queue length, and start I/O rate. The values are data elements DVAPCACT, DVAAVQUE, and DVASIO/DURATION.
QUEUE:	The estimated average internal queuing delay before IOS is able to start an I/O to the device. This value is estimated by subtracting the calculated values of control unit, device, and SIO time from the total estimated response time for the device.
CNTRL:	The estimated average control unit delay to obtain a free path to the device. This value is estimated based on percent control unit delay and I/O rate. These values are data elements DVAPCCUB and DVASIO/DURATION.
DEVICE:	The estimated average delay to reach the device when the control unit and channel are both free. This value is estimated based on percent reserve delay and the I/O rate. These values are data elements DVAPCRDL and DVASIO/DURATION.
SIO:	The estimated average device I/O time. This value includes channel control word (CCW) fetch, latency, data transfer, and missed rotational positioning sensing (RPS) reconnects. This

value is estimated based on the percent device activity and the I/O rate. These values are data elements DVAPCACT and DVASIO/DURATION.

To assist in the interpretation of the components of I/O, the values of the I/O component times are plotted over a range of 60 milliseconds. This plot is presented on the right-hand side of the report. The characters used to denote the I/O components are:

- Q - The internal queuing delay for the device.
- C - The control unit delay for the device.
- D - The device delay. Device delay results from RESERVE CCWs that are issued to the device by applications that need the device exclusively during some update process.
- S - The I/O time. As was mentioned above, this value includes CCW fetch, latency, data transfer, and missed RPS reconnects.

Each character represents one millisecond. The calculated values are rounded up by adding 0.49 milliseconds to the calculated values. Values exceeding 60 milliseconds are represented by a plus sign at the extreme right hand of the plot.

CA MICS Performance Manager																
DASD I/O RESPONSE TIME COMPONENT ANALYSIS																
INTERVAL START: 04JANyy:08:59																
INTERVAL END : 04JANyy:10:58																
DURATION : 2:00:00																
SYSID : SYSA																
-----RESPONSE TIME COMPONENTS (MS)-----																
DEV	DEVICE	SIO	%	-----AVERAGE TIME (MS)-----				1 2 3 4 5 6								
ADR	VOLSER	TYPE	RATE	BUSY	RESP	QUEUE	CNTRL	DEVICE	SIO	12345678901234567890123456789012345678901234567890						
240	SYST50	3330-11	3.0	9.8	42.3	5.6	2.1	2.2	32.3	Q	Q	Q	Q	Q	Q	Q
248	STOR11	3350	1.0	3.0	34.0	0.9	2.1	1.6	29.4	Q	C	D	S	S	S	S
25E	TTF569	3350	1.2	3.3	32.9	1.2	3.6	0.8	27.3	Q	C	C	D	S	S	S
26C	SP00L3	3380	2.8	6.3	30.3	4.3	1.2	2.1	22.8	Q	Q	Q	C	D	S	S
26D	SP00L4	3380	2.8	6.1	30.0	4.8	1.3	2.0	22.0	Q	Q	Q	C	D	S	S
260	SPCKPT	3380	3.9	14.9	64.9	1.5	0.0	25.1	38.4	Q	D	D	D	D	D	D
264	SP00L2	3380	3.2	11.1	50.2	9.3	3.2	3.4	34.2	Q	Q	Q	Q	Q	Q	Q
268	SP00L1	3380	2.6	6.3	32.9	5.3	1.3	2.4	23.8	Q	Q	Q	C	D	S	S
3E0	STOR13	3350	1.8	4.5	29.1	1.8	1.0	1.6	24.7	Q	C	D	S	S	S	S
3E1	PAG3E1	3350	2.3	13.3	73.0	8.3	5.2	1.3	58.3	Q	Q	Q	Q	Q	Q	Q
3E8	STOR14	3350	1.9	4.2	24.3	1.2	0.5	0.7	21.9	Q	D	S	S	S	S	S
322	AVPTW0	3330-11	1.1	4.1	38.4	0.3	0.0	0.0	38.2	S	S	S	S	S	S	S
328	LIBTS0	3330-11	1.2	2.5	22.1	0.7	0.4	0.6	20.5	Q	D	S	S	S	S	S
363	SFAD04	3350	1.1	2.4	22.8	0.6	0.5	0.0	21.7	Q	S	S	S	S	S	S
4B0	STORE5	3350	1.0	3.0	35.5	1.1	2.1	3.1	29.2	Q	C	D	D	S	S	S
467	TTF550	3350	1.1	1.9	20.1	1.7	1.1	0.2	17.0	Q	C	S	S	S	S	S
47A	TTF568	3350	1.1	3.2	32.6	1.7	1.0	1.6	28.3	Q	C	D	S	S	S	S
48E	TTF571	3350	1.2	2.2	21.5	1.3	0.8	0.8	18.6	Q	C	D	S	S	S	S
490	PAG490	3350	2.3	13.0	68.3	7.4	4.7	0.0	56.3	Q	Q	Q	Q	Q	Q	Q
52A	SYST55	3350	10.8	13.5	23.1	5.0	0.3	5.3	12.5	Q	Q	Q	Q	D	S	S
528	PAG528	3350	2.2	15.2	90.5	12.4	7.8	2.4	67.9	Q	Q	Q	Q	Q	Q	Q
53E	LIBSHR	3350	3.1	7.4	44.8	5.8	1.0	14.0	23.9	Q	Q	Q	Q	C	D	S
533	TTF533	3350	1.7	3.7	35.3	4.0	0.9	8.2	22.1	Q	Q	Q	C	D	S	S
535	XXXACC	3350	1.5	4.1	34.2	1.4	0.9	5.0	27.0	Q	C	D	D	S	S	S
9C0	CATPAK	3330	3.9	8.5	25.7	3.2	0.7	0.0	21.8	Q	Q	C	S	S	S	S
9C9	PAG9C9	3350	2.3	12.9	67.5	7.5	4.6	0.0	55.5	Q	Q	Q	Q	Q	C	S
9D8	TTF524	3350	1.0	1.7	18.9	1.7	0.0	0.0	17.2	Q	Q	S	S	S	S	S
93D	SFAD01	3350	1.5	3.1	23.0	1.9	0.5	0.3	20.4	Q	S	S	S	S	S	S
961	SFAD07	3350	2.8	5.3	22.5	1.9	0.4	1.1	19.1	Q	D	S	S	S	S	S
963	TTF559	3350	2.8	5.4	29.2	4.1	1.1	4.7	19.3	Q	Q	Q	C	D	S	S
965	PACK02	3350	1.6	4.0	34.6	2.9	2.5	3.1	26.0	Q	Q	C	C	D	S	S
966	XXXACT	3350	1.9	4.4	25.6	1.4	0.1	0.9	23.3	Q	D	S	S	S	S	S
967	PAG967	3350	1.8	3.8	25.5	1.9	1.3	1.3	21.0	Q	C	D	S	S	S	S

Figure 2-97. I/O Response Time Component Analysis Report for MVS/370

2.4.4.2 Technique Tutorial

The I/O response time is the most significant part of response or turnaround time for all but the most CPU bound applications. Over the past decade, we have witnessed CPU internal processing speeds increase 10 to 20 times. However, the average response time of a 3390 is only about two and one-half times faster than the 2314, which was initially introduced with the IBM S/360 series of processors. These trends in device and processor speeds only serve to exacerbate the performance problems that result from I/O delays.

MVS/XA AND MVS/ESA

Under MVS/XA and MVS/ESA, the components of I/O time are physically measured by the hardware. I/O response time is divided into the following components:

- o Queue time - The time that the I/O was queued waiting for the path or the device.
- o Pending time - The time after the I/O was issued, before it initially connected to the device.
- o Disconnect time - The time that the device was busy with seek, latency, and missed RPS rotations.
- o Connect time - The time that the device was connected to the path in data transfer mode.
- o ESCON port busy time - The time that the device was delayed due to an ES Connection Director (ESCD) port being busy.

MVS/370

In environments prior to MVS/XA, the I/O response time as perceived by the user application or transactions can be calculated as the sum of a number of components. By breaking response time into components, you can identify problems that can be addressed individually in order to improve performance. Two analytic methods have been published for dividing I/O response into component parts:

- o A technique that was published by Dorn, Friesenborg, and Hallsworth that breaks I/O response time into four components.

- o A technique presented by Sprinzen that breaks I/O response time into eight components.

Although the second technique is more detailed, it requires that a number of assumptions be satisfied and that detailed knowledge be available about the switching and composition of I/O strings. RMF data do not satisfy the necessary requirements for this technique. Hence, I/O component analysis uses the first method. Neither of these two analytic methods is exact; however, they both provide valuable insight into the delays encountered by DASD SIOs.

The interpretation of these values is discussed in the following sections:

- 1 - I/O Component Analysis for MVS/XA and MVS/ESA
- 2 - I/O Component Analysis for MVS/370

2.4.4.2.1 I/O Component Analysis for MVS/XA and MVS/ESA

I/O component analysis for MVS/XA is greatly simplified, since four components of I/O service time are directly measured and need not be estimated with queuing relationships, as was the case with pre-MVS/XA systems. The four components of I/O service time measured by MVS/XA systems are:

- o Queue - The average amount of time that I/O requests to this device were held enqueued within IOS before a SSCH instruction could be issued. Such an enqueue occurs when this particular system already has an I/O operation in progress on a device at the time another I/O operation is requested by this same system.
- o Pending - The time that the I/O was blocked by other I/Os that used some part of the physical path to the device. This can be caused by head of string or control unit contention or reserve delays issued from other systems.
- o Disconnect - The time that the device spent in seek, device latency, and extra rotations due to rotational positioning sensing (RPS) misses.
- o Connect - The time that the device was connected to the path. This includes the time required for both channel control word (CCW) and data transfer.
- o ESCON - The time that the device was delayed due to an ES Connection Director (ESCD) port being busy.

The first three of these components can be addressed by tuning. The fourth component is a function of data set block sizes and can only be addressed by changing the characteristics of the data sets on the device. The last of these components can be managed by adjusting the I/O load on the ES Connection Director ports.

The queue component of I/O response time is the sum of the time that the I/Os wait on the device's UCB (that is, device busy condition) and the time that the I/O queues behind the LCU (logical control unit) due to the utilization of the path by other I/Os. If the utilization of the device is in excess of 30 percent, the I/O queue time can be reduced by spreading the data sets on the volume across two or more devices. If the device utilization is low, the only way to address I/O queuing is to move the volume to a lower activity string.

The pending component of I/O response time represents the average time between successfully issuing the start subcommand (SSCH) is issued by the logical control unit (LCU) and acceptance of the first command by the the device. Two factors contribute to this delay. They are the utilization of the path (for example, channels or control units) by other I/Os from the system being measured, or device contention conditions that result from I/Os issued by other systems. These effects may be evaluated using the DASD Skew and Shared DASD Analysis Reports.

The disconnect component of I/O response time represents the portion of the device service time during which it is disconnected from the channel. This time includes seek, latency, and extra revolutions that result from RPS missed reconnects. The value shown in the report should be compared to the ideal disconnect time for the device type.

The ideal disconnect time may be calculated by dividing the sum of the average seek time and rotational delay for the device type by two. For example, the average seek time of a 3350 is 25 milliseconds and the rotational delay is 16.7 milliseconds. Hence, the ideal disconnect time for 3350 type devices is approximately 21 milliseconds.

By comparing the ideal disconnect time to the actual for the device, you can determine the impact of RPS missed reconnects. If this time is significant, you can move the volume to a less active device string where the RPS miss contribution is lower. You can use the DASD Skew Report to evaluate strings where RPS misses are causing performance problems.

Delays due to ES Connection Director port busy occur when I/O actions are delayed by the temporary unavailability of an ESCON director port. Unless a dedicated port-to-port connection from a channel to a controller has been defined within the ESCON director, a port-to-port connection will be established dynamically only for the duration of an I/O operation, with the I/O processor requesting termination of the connection upon I/O completion. A channel along which the I/O process is initiated must pass a request that the ESCON director establish a dynamic connection to any of the defined routes to the device for which the I/O operation is intended. Each controller attached to the device is attached to a port on the ESCON director. If each of these ports is busy, an ESCON port busy condition will occur.

Of course, if there is direct channel connection to a controller, ESCON port delays are not possible.

It is also possible to link two ESCON directors together, for instance, in order to increase the distance permitted between processors and devices. If this has been done, the link between the directors may also be a source of contention for which the port busy condition will occur.

2.4.4.2.2 I/O Component Analysis for MVS/370

DASD I/O response time is the sum of a number of unique events, as is shown in the diagram in Figure 2-98. The actual events are shown above the time line, and the calculated components of I/O response time are shown below. (Note that the asterisks on the time line indicated data transfer.) The first technique, published by IBM for MVS/370 systems, breaks I/O response time into four distinct components:

- o Queue - The sum of the time IOS spent queuing for the device and the logical channel.
- o Control Unit - The time that the channel was available, but the control unit(s) for the device were busy.
- o Device - The delay to reach the device caused by reserves from other systems or head-of-string contention.
- o SIO - The time that the I/O was active on the device. This value includes seek, set sector, latency, RPS reconnect, and data transfer.

The values for each of these four components can be approximated from standard RMF measurements.




Figure 2-98. I/O Response Time Components

The four components of I/O response time are computed from six standard CA MICS data elements:

- o DVAPCACT - Device percent busy
- o DVAPCRDL - Device percent reserve delay
- o DVAPCCUB - Percent control unit busy
- o DVAAVQUE - Device average queue length
- o DVASIO - The total number of SIOs serviced by the device
- o DURATION - The length of the interval

For additional information on these values, consult the CA MICS Data Dictionary.

The total number of start I/Os and the duration are used to compute the start I/O rate:

$$\text{SIORATE} = \text{DVASIO} / \text{DURATION} \quad (\text{Equation 1})$$

The SIORATE value will be used in equations in the following paragraphs.

SIO time is computed based on the percent device busy value and the start I/O rate. This calculation is shown below:

$$\text{SIO} = (\text{DVAPCACT}/100) / \text{SIORATE} \quad (\text{Equation 2})$$

The most significant portion of SIO time is device rotational delay. For devices on busy paths or strings, the path to the CPU can be busy when the device attempts to reconnect to the channel. This causes the reconnection to fail, and the I/O must then wait a complete rotation (16.7 milliseconds for most device types) before another reconnect can be attempted. RPS reconnect probabilities are discussed in the DASD skew technique tutorial (Section 2.4.5.2).

Device time is computed based on the percent reserve delay and the start I/O rate. This calculation is shown below:

$$\text{DEVICE} = (\text{DVAPCRDL}/100) / \text{SIORATE} \quad (\text{Equation 3})$$

Device time, that is, the percent of time the system could not access the device even though the UCB (unit control block) indicated that it was free and the channel path and control

unit were free, is caused by three problems:

- o Shared DASD contention
- o Reserves issued by another system
- o Head-of-string contention

The first two of these problems can be investigated using the Shared DASD Analysis Report. The DASD Skew Analysis Report can provide insight into head-of-string contention problems.

Control unit delay occurs when the channel is free, but the control unit is busy supporting an I/O to another device or string. Control unit delay is calculated based on percent control unit delay and the start I/O rate. This calculation is shown below:

$$\text{CNTLUNIT} = (\text{DVAPCCUB}/100) / \text{SIORATE} \quad (\text{Equation 4})$$

In most cases, this time is very small, that is, less than five milliseconds.

Unfortunately, the internal queuing time cannot be directly estimated. However, we can exploit the fact that we have already computed the other three components needed to derive total response time. The total response time is computed based on the device busy value, the average queue length, and the start I/O rate. This calculation is shown below:

$$\text{TOTAL} = (\text{SIO} + \text{DVA AVQUE}) / \text{SIORATE} \quad (\text{Equation 5})$$

Knowing the total time, we can then calculate the internal queuing delay, as shown below:

$$\text{QUEUE} = \text{TOTAL} - \text{DEVICE} - \text{CNTLUNIT} - \text{SIO} \quad (\text{Equation 6})$$

High values of internal queuing are caused by two problems:

- o Very high I/O request rates to the device. To solve this problem, you should migrate data sets to other devices to balance I/O activity. Most often, this problem results from an application putting multiple high activity data sets on a single device. You can identify the performance group using a device by investigating the EXCP counts in the BATWDA File.
- o Significant queues for the logical channel. You can determine this problem by investigating the average queue length value of the standard RMF logical channel report that may be produced using CA MICS' standard RMF

reports compendium program. If this value is high, you may have to reconfigure your logical and physical channels.

2.4.4.3 Case Study

The I/O service time as measured by the system is the sum of a number of unique components. Each of these components can be addressed by different tuning techniques to minimize the overall response time on the device. As we discussed in Section 2.4.4.2, Technique Tutorial, we can use some simple relationships to approximate the values of each of the response time components. The DASD I/O Component Analysis Report allows you to apply this methodology to analyze the individual device response time characteristics of your systems.

For this case study, we used the DASD I/O Component Analysis Report to analyze the response times of devices on an IBM 3084 with a SYSID of SYSA running MVS/XA. The peak morning periods, hours 9 and 10, were analyzed for the previous day. The CA MICS database unit that contained the installation's hardware data was denoted by database identifier B. The inquiry panel for the study is shown in Figure 2-99.

```
/----- IOA - DASD I/O Component Analysis -----\
| Command ==>                                         |
| Composing CA MICS Inquiry: IOA18 - DASD Component Case Study |
| Time-span ==> DETAIL (DETAIL/DAYS)                  Cycle(s) ==> 01 - 01 |
| Data base ==> B (PRIMARY) _                         _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
| Report by time interval ==> N (Y/N) Minimum I/O rate ==> 1 (.001 - 9999) |
| Usage qualification processing:                     |
| CA MICS element to test==> _____ (Any numeric HARDVA element) |
| Operator to employ ==> ____ (LE,LT,GT,GE) |
| Numeric value to test ==> _____ (.00000000 - 99999999) |
| Data Selection:                                     Allow run-time override |
| Date ==> _____ - _____ - _____ ==> N (Y/N) |
| Hour ==> 09 - 10 ____ - ____ - ____ ==> N (Y/N) |
| Zone ==> _ _ _ _ _ ==> N (Y/N) |
| Sysid ==> SYSA _ _ _ _ _ ==> N (Y/N) |
| Specify extended options ==> N (Y/N) |
\-----/
```


Figure 2-99. Case Study - DASD I/O Component Analysis Inquiry Panel
The DASD I/O Component Analysis Report produced by the panel shown above is provided in Figure 2-100.

You should carefully examine a number of values on the report, for example:

- o The queue time should be proportional to the utilization of the device. Utilizations in excess of 30 to 35 percent should be avoided, except for paging or swapping device (see Section 2.3, ASM Analysis). High utilization will result in significant internal queuing for the device.
- o Pending delay resulting from device reserves that are issued by applications on other systems or controller contention delays should be small, except for very specialized volumes like a checkpoint volume.
- o Large disconnect delays may indicate heavy seek activity or substantial RPS misses. With multiple paths to a device, the possibility for RPS miss is substantially reduced. Under these conditions, heavy seek activity is a likely culprit. In this particular case study, the 3380 devices were the standard models that can and did experience substantial internal string contention, leading to substantially more disconnect time than the 3390 devices, which were operating as DSLE devices with four available paths.
- o Connect time is almost completely dependent upon the amount of data being transferred, and as such, is time spent doing useful work.
- o Delays due to ES Connection Director port busy occur when I/O actions are delayed by the temporary unavailability of an ESCON director port. Large ESCON port busy delays may indicate contention for a controller attached to an ESCON director or heavy use of director-to-director links in a multiple director connection.

By using the DASD I/O Component Analysis Report, you can identify devices that are causing significant performance problems.

2.4.4.4.1 DASD I/O Component Analysis Inquiry Panel

Figure 2-101 shows the DASD I/O Component Analysis inquiry panel.

```

/----- IOA - DASD I/O Component Analysis -----\
| Command ==>                                     |
|                                                    |
| Composing CA MICS Inquiry: IOA20 - DASD Component Analysis |
|                                                    |
| Time-span ==> DETAIL (DETAIL/DAYS)                Cycle(s) ==> 01 - 02 |
| Data base ==> A (PRIMARY) _                      _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
|                                                    |
| Report by time interval ==> N (Y/N) Minimum I/O rate ==> 1 (.001 - 9999) |
|                                                    |
| Usage qualification processing:                    |
| CA MICS Element to test==> _____ (Any numeric HARDVA element) |
| Operator to employ ==> ____ (LE,LT,GT,GE) |
| Numeric value to test ==> _____ (.00000000 - 99999999) |
|                                                    |
| Data Selection:                                     Allow run-time override |
| Date ==> _____ - _____ - _____ ==> N (Y/N) |
| Hour ==> ____ - ____ - ____ - ____ ==> N (Y/N) |
| Zone ==> _____ ==> N (Y/N) |
| Sysid ==> _____ ==> N (Y/N) |
|                                                    |
| Specify extended options ==> N (Y/N) |
|                                                    |
\-----/

```

Figure 2-101. DASD I/O Component Analysis Report Inquiry Panel

The following is a discussion of the fields on this panel.

Timespan: The specification of whether the data is to be taken from the DETAIL or DAYS timespan.

Cycle(s): One or more cycles of the CA MICS files that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a cycle from the CA MICS file usually contains a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided

by the DATE and HOUR control parameters.

Database:	The single-character CA MICS identifiers for the databases you want to access.
Allow Runtime Override:	The indication of whether the database identifier and/or the Cycles specification can be overridden at inquiry execution time.
Report by time interval:	The specification of whether a report should be generated for each interval specified or summarized into a single report.
Minimum I/O rate:	A filter for the exclusion of low I/O activity volumes. It specifies the minimum per second I/O rate for volume inclusion in the analysis. This value may range between 0.001 and 999. This allows you to exclude devices that were either unused during the interval or received very little use. If you use this control parameter, the I/O filter applies to all I/O component and shared-DASD reports.
CA MICS Element to test:	In addition to the Minimum I/O rate, the Shared DASD I/O Component Analysis Report inquiry can test any other numeric data element in the HARDVA File for a specific condition. This can be used to further subset the data to be analyzed.
Operator to employ:	The boolean operator to employ - LE, LT, GT, or GE.
Numeric value to test:	The value against which the data element is tested.
Date:	The day(s) to be used for analysis. This field is normally used with the IOA Shared DASD Analysis inquiry to drop the extra day from a cycle that is due to the timing of the SMF dump.
Hour:	An hour or range of hours for which data is reported.
Zone:	The specification that zone selection is to be applied. If zone 1 definition is 8:00 a.m. to 6:00 p.m., daily reports covering that time period can be specified by using the ZONE

field set to 1.

Sysid: The SYSID of the data to be reported.

Specify extended options: The specification that allows you to supply additional report options. These options include special SAS data selection statements, non-standard input, and output files. These options are discussed in the sections that follow.

The operation of the time selection control parameters depends on definitions made during the installation of CA MICS. If you are not sure about the zone definitions for your installation or the CA MICS database update procedures, consult your CA MICS System Administrator.

2.4.4.4.2 Extended Options Menu Panel

If you enter Y (Yes) in the Specify extended options field, the Extended Options Menu, shown in Figure 2-102, is displayed.

```
/----- Extended Options Menu -----\
| Command ==>
|
| Inquiry Step: IOA DASD Component Analysis
|
| Specify global data selection exit          ==> N (Y/N/R)
|
| Specify alternate input file                ==> N (Y/N/R)
| Save selected input observations            ==> N (Y/N/R)
|
|
|
|
|
|
```

Figure 2-102. DASD I/O Component Analysis Report Extended Options Menu Panel

This Extended Options Menu panel allows you to specify that:

- o Global data selection statements are to be defined.
- o Non-CA MICS input files are to be defined.
- o Non-CA MICS output files are to be defined.

The default for all three options is N (No). The R (Reset) option allows you to return the option to its original contents. For detailed information about the extended options, press the help key (usually PF1) when you are at the Extended Options Menu.




Figure 2-103. DASD I/O Component Analysis Report Global Data Selection Panel

2.4.4.4.4 Alternate Input File Panel

If you enter Y (Yes) in the Specify alternate input file field in the Extended Options Menu, the Alternate Input File panel, shown in Figure 2-104, is displayed.

```
/----- Alternate Input File -----  
| Command ==>  
|  
| Inquiry Step: IOA DASD Component Analysis  
|  
| Specified Time-span: DETAIL  
|  
| Equivalent CA MICS file(s): HARDVA  
| Alternate SAS file(s)    ==> _____  
|  
| SAS data set allocation instructions:  
|   Data set name          ==> _____  
|   Disposition            ==> SHR  (SHR/OLD)  
  
|  
|  
|  
|  
|
```

Figure 2-104. DASD I/O Component Analysis Report Alternate Input File Panel

This panel allows you to specify a SAS file that is not a part of the standard CA MICS online database as input to the IOA Analysis routines. You must specify one or more indicated SAS file names and an MVS data set name. If you do not enclose the data set name in single quotes, your TSO prefix (or user ID) is added as the first portion of the data set name.

2.4.4.4.5 Save Selected Input Observations Panel

If you enter Y (Yes) in the Save selected input observations field in the Extended Options Menu, the Save Selected Input Observations Panel, shown in Figure 2-105, is displayed.

[illegible]

Figure 2-105. DASD I/O Component Analysis Report Save Selected Input Observations Panel

This panel allows you to specify a SAS file that is not a part of the CA MICS online database as output from the IOA Analysis routines. Use this option if you want to save a copy of the data that went into the report. You must enter one or more SAS file names for the equivalent CA MICS file(s) and the MVS data set name that will contain the SAS file(s). You can read the same data into a later inquiry by specifying the SAS file and data set names on the Specify alternate input file extended option.

2.4.5 DASD Skew Analysis Inquiry

Prior to the introduction of controllers with RPS avoidance, such as EMC and 3990-6 controllers, a condition existed called DASD skew. DASD skew occurs when one or two of the devices on a string of DASD are particularly busy and are adversely affecting the performance of the remainder of the devices on the string.

DASD skew is discussed in the technique tutorial (Section 2.4.5.2). The DASD Skew Analysis Report allows you to identify strings that have significant balance problems on their devices.

This section discusses the following topics:

- 1 - Standard Output
- 2 - Technique Tutorial
- 3 - Case Study
- 4 - Component Operation

2.4.5.1 Standard Output

Figure 2-106 shows the DASD Skew Analysis Report inquiry panel. A description of the inquiry fields is presented in Section 2.4.5.4.1 in the discussion on component operation.

```

/----- IOA DASD Skew Analysis Report -----\
| Command ==>                                     |
|                                                    |
| Composing CA MICS Inquiry: IOA25 - DASD Skew 10:00 / SYSA |
|                                                    |
| Time-span ==> DETAIL (DETAIL/DAYS)                Cycle(s) ==> 01 - 01 |
| Data base ==> B (PRIMARY) _                        _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
|                                                    |
| Report by time interval ==> N (Y/N)                |
|                                                    |
| Data Selection:                                     Allow run-time override |
|   Date          ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
|   Hour          ==> 10 - _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
|   Zone          ==> _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
|   Sysid         ==> SYSA _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
|                                                    |
| Specify extended options ==> N (Y/N)                |
|                                                    |
\-----/

```

Figure 2-106. IOA DASD Skew Analysis Report Inquiry Panel

Figure 2-107 shows a sample DASD Skew Analysis Report created by the inquiry entries shown above. The following items are shown in the figure:

INTERVAL START:	The time and date of the first HARDVA File observation included in the analysis.
INTERVAL END:	The time and date of the last HARDVA File observation included in the analysis.
DURATION:	The total number of hours summarized by report, rounded up to the nearest hour.
SYSID:	The system identifier.

CU ADR/ CU/PATH:	The control unit address, CA MICS data element CUADDR. For MVS/XA and ESA this column is replaced by the value CU/PATH. In MVS/XA or ESA control unit number is the equivalent data element to control unit address in MVS/370. For data derived from RMF 4.2 and above, the control unit name is also shown in the line just below the control unit number.
LOG. CHAN/ LCU:	The logical channel number that serves the control unit, CA MICS data element LOGCHAN. For MVS/XA and ESA, this column is replaced by the value LCU (logical control unit). The logical control unit data element was introduced in MVS/XA. There is no corresponding data element to logical channel in MVS/XA.
AV CHAN SERV MS/ AV LCU SERV MS:	The weighted average physical channel service time for the physical channels that comprise the logical channel. For MVS/XA and ESA, this column is replaced by the AV LCU SERV MS, the average connect time for devices on the logical control unit.
TOTAL SIO RATE/ TOTAL SSCH RTE:	The total rate of start I/Os (SIOs) issued per second to the devices on the string by the system. For MVS/XA and ESA, this column is replaced by TOTAL SSCH RTE, the total rate of start sub-channel commands issued per second to the string.
TOTAL % HOS:	The total utilization of the head of string (HOS). For pre-MVS/XA systems, the value is calculated by multiplying the total number of SIOs times the average channel service time in milliseconds, divided by the duration. For MVS/XA and ESA systems, the value is calculated by dividing the sum of the device connect times by the duration.

TYPE LINE: There are seven lines of detailed data presented for each device string. The TYPE LINE values describe the values that are found in the next eight columns, DEV 1 through DEV 8. These columns correspond to device addresses 0-7 and 8-F for non-3380 device types. For 3380s without the Device Level Selection feature, the first four columns are used to represent the internal device strings. They are:

String	Addresses
-----	-----
1	0,1,8,9
2	2,3,A,B
3	4,5,C,D
4	6,7,E,F

The types of lines are:

- o DEVADDR - The device address on the string, DEVADDR. For MVS/XA and ESA, this row is replaced by the value DEV NO., the corresponding data element to device address in MVS/370.
- o SIO RT - The rate of start I/Os (SIOs per second) issued to each device by the system, DVASIO. For MVS/XA and ESA systems, this row is replaced by SSCH RT, the rate of start subchannel commands issued to the device.
- o AV SERV - The average service time for the device, DVAVSTM.
- o % PATH - The percent of the path used by the device, computed like TOTAL % HOS. See above.
- o VOLSER - The volume serial number of the volume mounted on the device.
- o STGGRP - The DFSMS storage group name. DFSMS (Data Facility Storage Management Subsystem) allows for pooling physical DASD volumes. Each pool is defined by the installation's DASD administrator, and is managed by DFSMS.
- o DEVTYPE - The DASD device type of each reported volume. (For example, 3380J, or

3390).

- CV: The coefficient of variation of the non-zero percent path busy values for the devices on the string. It is important to note that idle devices are not included in the calculations. To calculate the coefficient of variation, at least one of the percent device busy values must be non-zero. The higher the coefficient of variation, the more significant the skew problem of the string.
- SKEW: The skew coefficient of the non-zero percent path busy values for the devices on the string. It is important to note that idle devices are not included in the calculations. To calculate the skew coefficient, at least two of the percent device busy values must be non-zero. The higher the skew coefficient, the more significant the skew problem of the string.

If the Device Level Selection (DLS) feature is in effect, then there will only be one set of values for CU/PATH, LCU, AV CHAN SERV MS, TOTAL SSCH RTE, TOTAL % HOS, CV, and SKEW presented for the up to 16 device addresses associated with each control unit or control unit pair.

If the Device Level Selection Enhanced (DLSE) feature is in effect, reflecting the utilization of 3990-2 or 3990-3 controllers, then there will only be one set of values for CU/PATH, LCU, AV CHAN SERV MS, TOTAL SSCH RTE, TOTAL % HOS, CV, and SKEW presented for the up to 64 device addresses associated with each logical control unit.

CA MICS Performance Manager																
DASD SKEW ANALYSIS																
INTERVAL START: 07SEPy:09:59																
INTERVAL END : 07SEPy:10:59																
DURATION : 2:00:00																
SYSID : XA01																
CU /	AV	LCU	TOTAL	TOTAL	TYPE										CV /	
PATH	LCU	SERV	MS	SSCH	RTE	% HOS	LINE	DEV 1	DEV 2	DEV 3	DEV 4	DEV 5	DEV 6	DEV 7	DEV 8	SKEW
570-3 3880	0009	25.06		5.7	14.26		DEV NO.	574	575	57C	57D					1.69
							SSCH RT	5.3	0.2	0.2	0.0					1.73
							AV SERV	37.37	17.63	27.38	0.00					
							% PATH	14.02	0.12	0.12	0.00					
							VOLSER	TS0920	TS0900	TS0904	VSAM04					
							STGGRP	TS0DEV	TS0PRD	TS0DEV	DBMS02					
							DEVTYPE	3380	3380	3380	3380					
8C0-0 3390-2	0018	3.99		14.3	6.13		DEV NO.	8C0	8C1	8C2	8C3		8C5	8C6	8C7	--SEE-- * DLSE*
							SSCH RT	0.0	0.0	0.0	0.0		4.0	1.3	0.1	-BELOW- ANY 4
							AV SERV	0.78	0.77	0.77	0.77		9.41	4.50	11.65	DEVICES
							% PATH	0.00	0.00	0.00	0.00		1.14	0.33	0.04	CAN BE
							VOLSER	M24RSS	M32SPL	M32RSS	M28RSS		M18SPL	M18RS1	TAS002	ACTIVE
							STGGRP	N/A	SYSTEM	N/A	N/A		SYSTEM	N/A	TECSUP	
							DEVTYPE	3390	3390	3390	3390		3390	3390	3390	
8C0-0 3390-2	0018	-----SEE ABOVE-----					DEV NO.	8C8	8C9	8CA	8CB	8CC	8CD	8CE	8CF	--SEE-- * DLSE*
							SSCH RT	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	-BELOW- ANY 4
							AV SERV	0.78	0.85	4.89	8.25	0.79	7.16	8.52	8.45	DEVICES
							% PATH	0.00	0.00	0.00	0.02	0.00	0.00	0.00	0.06	CAN BE
							VOLSER	HDK008	MSH009	MSH010	MSH011	MSH012	MSH013	TAS020	TS0021	ACTIVE
							STGGRP	N/A	MSHGRP	MSHGRP	MSHGRP	MSHGRP	MSHGRP	TECSUP	TS0DEV	
							DEVTYPE	3390	3390	3390	3390	3390	3390	3390	3390	
8C0-0 3390-2	0018	-----SEE ABOVE-----					DEV NO.	8D0	8D1	8D2	8D3	8D4	8D5	8D6	8D7	--SEE-- * DLSE*
							SSCH RT	0.2	0.1	0.6	0.0	0.0	0.0	0.0	0.0	-BELOW- ANY 4
							AV SERV	15.74	8.45	7.08	8.61	3.40	0.78	0.81	0.83	DEVICES
							% PATH	0.22	0.03	0.15	0.00	0.00	0.00	0.00	0.00	CAN BE
							VOLSER	TS0029	TS0026	TS0027	TS0028	TS0030	TS0031	TS0032	TS0033	ACTIVE
							STGGRP	TSUPRD	TS0DEV	TS0DEV	TS0PRD	N/A	N/A	N/A	N/A	
							DEVTYPE	3390	3390	3390	3390	3390	3390	3390	3390	
8C0-0 3390-2	0018	-----SEE ABOVE-----					DEV NO.	8D8	8D9	8DA	8DB	8DC	8DD	8DE	8DF	2.66 * DLSE*
							SSCH RT	0.0	0.0	0.2	0.1	0.0	0.3	0.0	0.0	0.97 ANY 4
							AV SERV	4.59	5.72	9.87	7.50	0.77	13.90	0.80	14.23	DEVICES
							% PATH	0.00	0.00	0.09	0.04	0.00	0.12	0.00	0.02	CAN BE
							VOLSER	TS0034	TS0035	TS0036	TS0037	TEC003	TMP072	MSH015	MSH014	ACTIVE
							STGGRP	TSOPRD	TSOPRD	TS0DEV	N/A	TECSUP	TECSUP	MSHGRP	MSHGRP	
							DEVTYPE	3390	3390	3390	3390	3390	3390	3390	3390	

Figure 2-107. IOA DASD Skew Analysis Report for MVS/XA

2.4.5.2 Technique Tutorial

DASD skew is a significant problem because it allows one busy device on a string to adversely affect the performance of the remainder of the devices on the string. To understand the nature of the problem, it is first necessary to understand the concepts of rotational position sensing (RPS) devices and RPS miss probabilities.

With the introduction of the 3330 device in the early 1970s, IBM changed the channel protocol for disks from a selector-like channel protocol, where the channel was left busy for much of the time while the disk arm moved (seek) and while the disk read to the desired physical record, to a rotational position sensing protocol. Under the RPS protocol, the device disconnects from the channel while it positions itself to read the desired record. Since data transfer represented only a small percentage of the time that pre-RPS channels were busy, there was opportunity for significant performance improvements by allowing the channel to service other requests while the device moved to the desired record.

In theory, the device senses (via a value provided in a set sector channel control word (CCW)) that the head is about to encounter the desired record, and reconnects to the channel. Unfortunately, if the channel is busy when the device encounters the record, an RPS miss occurs and the device must wait a full rotation (16.7 milliseconds for most device types) until the head encounters the record again. It is easy to conclude that even an occasional RPS miss could significantly degrade the performance of a device, since an RPS miss introduces a delay equal to about one half of the average device service time for 3330, 3350, and 3380 device types.

Fortunately, the probability of an RPS miss is a simple function of the utilization of the channel that supports the device. You are probably familiar with the rule of thumb that says that block multiplexor channels should never be more than 30 percent busy. This rule of thumb can easily be understood by considering the equation for RPS miss probabilities. Though its derivation is beyond the scope of this tutorial, an understanding of the rule can be achieved by considering the following equation:

$$u p = \frac{p}{(1-u)} \quad \text{(Equation 7)}$$

where: p - is the probability of an RPS miss.

u - is the utilization of the path, i.e., channel.

The behavior of this equation is best demonstrated by considering the following table of path utilizations that vary from 5 to 50 percent. Consider:

u	p
5	0.05
10	0.11
15	0.18
20	0.25
25	0.33
30	0.43
35	0.54
40	0.67
45	0.82
50	1.00

This equation is an example of an exponential relationship. That is, very small changes in the independent variable u can have profound effects on the dependent variable p.

Now that you understand RPS reconnect probabilities, we can discuss the problems that are introduced by DASD skew. DASD skew describes a condition where one or two of the devices on a string are very heavily utilized. For example, consider the following table of eight disk devices and the percent of the path (i.e., channel) used by each of the devices.

DEVICE	% PATH (u)
1	3
2	12
3	1
4	0
5	2
6	0
7	3
8	1

To simplify the calculation of the RPS miss probabilities, we assume that the string of devices is on a dedicated channel. The probability of an RPS miss for device 2 is dependent on the total utilization of the path by the other devices on the string. (After thinking for a moment, it is clear that a

device cannot contend with itself for the channel. It is this insight that is key to the understanding the problem caused by DASD skew.) Therefore, the probability of an RPS miss for device 2 is a function of the sum of the path utilizations of the other devices, i.e., $3+1+0+2+0+3+1=10\%$. Thus, the RPS miss probability is:

$$p = \frac{.1}{(1-.1)} = 0.11$$

The impact of DASD skew can be most easily seen by calculating the RPS miss probability for a lightly loaded device like device 8. The probability of an RPS miss for device 8 is a function of the sum of the path utilizations of the other devices, i.e., $3+12+1+0+2+0+3=21\%$. Thus, the RPS miss probability is:

$$p = \frac{.21}{(1-.21)} = 0.27$$

From this simple exercise, we can conclude that the ideal string would have the path utilizations of all of its devices approximately equal.

One question that this argument introduces in an MVS/370 environment is, how can we calculate the utilization of the path for each device? The path utilization is calculated by multiplying the average service time of the logical channel that serves the device by the SIO rate for the device.

Although logical channel service time is not reported by standard RMF reports, it can be approximated easily for MVS/370 by using the service times of the physical channels that comprise the logical channel. For example, consider logical channel 7 that is comprised of physical channels 4 and 8. The average logical channel service time is computed by taking the weighted average (weighted by the number of SIOs processed by each of the physical channels) of the average physical channel service times for channels 4 and 8.

In MVS/XA this is not necessary, since the device connect time is reported by the hardware and the average value is maintained in CA MICS data element DVAVCNN.

Once the path utilizations of the devices on a string are known, statistical techniques can be applied to the values to statistically quantify the skewness of the devices on the

string. There are two statistical coefficients that indicate the skewness of a distribution:

- o The coefficient of variation, CV
- o The coefficient of skewness, SK

The coefficient of variation is computed from the mean and standard deviation of the distribution of the device path utilization values. The coefficient of variation is computed as:

$$CV = \frac{\text{standard deviation}}{\text{mean}} \quad (\text{Equation 8})$$

The larger the value of the coefficient of variation, the more significantly skewed the distribution.

The second statistical technique for evaluating the skewness of a distribution is the coefficient of skewness. The coefficient of skewness is computed from the mean, median, and standard deviation. The coefficient of skewness is computed as:

$$SK = \frac{3 * (\text{mean} - \text{median})}{\text{standard deviation}} \quad (\text{Equation 9})$$

The most important consideration when computing these coefficients is whether or not idle devices should be included in the computations. Clearly, a string with seven idle devices and one busy device will have very high coefficients of variation and skewness. However, there is no RPS miss problem, since the other devices are unused. Therefore, the algorithms implemented in the I/O Analysis routines exclude idle devices from the calculation of the coefficients.

Of the two coefficients for evaluating the statistical behavior of the devices on the string, the coefficient of skewness is probably a more reliable indicator. Also, you should note that the skewness calculations are made on a system-by-system basis, i.e., the utilization of the control unit and head of string (both parts of the I/O path) by other systems is not accounted for by the current implementation of the algorithm.

CA MICS Performance Manager														
DASD SKEW ANALYSIS - DAYS LEVEL														
13:08 WEDNESDAY, DECEMBER 19, yyyy														
INTERVAL START: 15NOVyy:09:52														
INTERVAL END : 15NOVyy:10:52														
DURATION : 1:00:05														
SYSID : SYSA														
CU /	AV	LCU	TOTAL	TOTAL	TYPE	DEV 1	DEV 2	DEV 3	DEV 4	DEV 5	DEV 6	DEV 7	DEV 8	CV /
PATH	LCU	SERV	MS	SSCH	RTE	% HOS	LINE							SKEW
1A0-0 0005	10.26		33.2	34.08	DEV NO.	1A0	1A1	1A2	1A3	1A4	1A5	1A6	1A7	--SEE-- * DLS *
					SSCH RT	0.3	0.0	0.0	0.0	0.0	0.0	0.0	0.3	-BELOW- ANY TWO
					AV SERV	30.63	0.00	37.06	0.00	0.00	0.00	0.00	21.60	DEVICES
					% PATH	0.12	0.00	0.00	0.00	0.02	0.00	0.00	0.21	CAN BE
					VOLSER	WORK21	XTG001	BENSC2	FYQ001	WORK30	WORK31	PIAN01	WMN017	ACTIVE
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J	
1A0-0 0005	-----SEE ABOVE-----				DEV NO.	1A8	1A9	1AA	1AB	1AC	1AD	1AE	1AF	2.14 * DLS *
					SSCH RT	0.0	0.0	0.0	31.5	0.0	0.1	0.0	1.1	1.36 ANY TWO
					AV SERV	0.00	0.00	0.00	22.16	0.00	10.21	0.00	22.48	DEVICES
					% PATH	0.00	0.00	0.00	32.96	0.00	0.02	0.00	0.76	CAN BE
					VOLSER	PIAN05	FAT1B9	CIS001	FAT1BB	FAT1BC	FAT1BD	FAT1BE	TECSYS	ACTIVE
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J	
1B0-0 0005	9.30		15.7	14.60	DEV NO.	1B0	1B1	1B2	1B3	1B4	1B5	1B6	1B7	--SEE-- * DLS *
					SSCH RT	.15	0.4	0.0	0.4	0.1	0.0	3.3	0.0	-BELOW- ANY TWO
					AV SERV	10.78	22.82	0.00	23.73	15.84	18.48	16.26	16.53	DEVICES
					% PATH	0.32	0.53	0.00	0.21	0.02	0.00	0.00	0.00	CAN BE
					VOLSER	WORK22	WORK26	WORK23	CH0100	PIAN02	PIAN04	BES003	BES001	ACTIVE
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J	
1B0-0 0005	-----SEE ABOVE-----				DEV NO.	1B8	1B9	1BA	1BB	1BC	1BD	1BE	1BF	2.25 * DLS *
					SSCH RT	10.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.13 ANY TWO
					AV SERV	24.30	0.00	48.81	0.00	0.00	0.00	0.00	0.00	DEVICES
					% PATH	12.71	0.00	0.02	0.00	0.00	0.00	0.00	0.00	CAN BE
					VOLSER	WMN003	WMN004	IPN001	WMN006	PIAN03	IPN002	CS0101	PIAN06	ACTIVE
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J	
1C0-0 0006	3.51		30.2	10.56	DEV NO.	1C0	1C1	1C2	1C3	1C4	1C5	1C6	1C7	--SEE-- * DLS *
					SSCH RT	0.0	0.0	0.1	0.1	0.5	8.5	2.1	0.6	-BELOW- ANY TWO
					AV SERV	15.81	0.00	38.00	27.63	15.84	18.48	16.26	16.53	DEVICES
					% PATH	0.00	0.00	0.11	0.08	0.10	3.30	0.41	0.15	CAN BE
					VOLSER	SPL007	SPL008	WMN025	WMN026	WJEP03	WJEP04	FEJ000	DLIB23	ACTIVE
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J	
1C0-0 0006	-----SEE ABOVE-----				DEV NO.	1C8	1C9	1CA	1CB	1CC	1CD	1CE	1CF	1.24 * DLS *
					SSCH RT	0.0	0.0	3.0	0.0	12.0	1.3	0.0	2.1	1.69 ANY TWO
					AV SERV	20.60	28.42	31.04	35.61	19.62	17.51	0.00	21.04	DEVICES
					% PATH	0.00	0.00	1.90	0.00	3.29	0.32	0.00	0.90	CAN BE
					VOLSER	TEC003	WJEP18	IPN004	MICS02	FAT25C	FAT25D	WMN012	WMN013	ACTIVE
					STGGRP	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	
					DEVTYPE	3380J	3380J	3380J	3380J	3380J	3380J	3380J	3380J	

Figure 2-109. Case Study - DASD Skew Analysis Report

2.4.5.4 Component Operation

The inquiry fields allow you to define the CA MICS database cycles to be used as input, the time periods to be analyzed, and the reports to be produced. Figures 2-110 through 2-114 show the inquiry panels used to produce the DASD Skew Analysis Report.

The following sections describe each control panel:

- 1 - DASD Skew Analysis Report Panel
- 2 - Extended Options Menu Panel
- 3 - Global Data Selection Panel
- 4 - Alternate Input File Panel
- 5 - Save Selected Input Observations Panel

2.4.5.4.1 DASD Skew Analysis Report Panel

Figure 2-110 shows the DASD Skew Analysis Report inquiry Panel.

```

/----- IOA DASD Skew Analysis Report -----\
| Command ==>                                     |
|                                                    |
| Composing CA MICS Inquiry:  IOA29  - Detail Level DASD Skew Analysis |
|                                                    |
| Time-span ==> DETAIL (DETAIL/DAYS)                Cycle(s) ==> 01 - 02 |
| Data base ==> A (PRIMARY) _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
|                                                    |
| Report by time interval ==> N (Y/N)                |
|                                                    |
| Data Selection:                                     Allow run-time override |
|   Date           ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
|   Hour           ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
|   Zone           ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
|   Sysid          ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
|                                                    |
| Specify extended options ==> N (Y/N)                |
|                                                    |
\-----/

```

Figure 2-110. IOA DASD Skew Analysis Report Inquiry Panel

The following is a discussion of the fields on this panel.

- Timespan:** The specification of whether the data is to be taken from the DETAIL or DAYS timespan.
- Cycle(s):** One or more cycles of the CA MICS files that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a cycle from the CA MICS file usually contains a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided by the DATE and HOUR control parameters.
- Database:** The single-character CA MICS identifiers for the databases you want to access.

Allow Runtime Override:	The indication of whether the database identifier and/or the Cycles specification can be overridden at inquiry execution time.
Report by time interval:	The specification of whether a report should be generated for each interval specified or summarized into a single report.
Date:	The day(s) to be used for analysis. This field is normally used with the IOA Analysis inquiry to drop the extra day from a cycle that is due to the timing of the SMF dump.
Hour:	An hour or range of hours for which data is reported.
Zone:	The specification that zone selection is to be applied. If zone 1 definition is 8:00 a.m. to 6:00 p.m., daily reports covering that time period can be specified by using the ZONE field set to 1.
Specify extended options:	The specification that allows you to supply additional report options. These options include special SAS data selection statements, non-standard input, and output files. These options are discussed in the sections that follow.

The operation of the time selection control parameters depends on definitions made during the installation of CA MICS. If you are not sure about the zone definitions for your installation or the CA MICS database update procedures, consult your CA MICS System Administrator.

2.4.5.4.2 Extended Options Menu Panel

If you enter Y (Yes) in the Specify extended options field, the Extended Options Menu, shown in Figure 2-111, is displayed.

```

----- Extended Options Menu -----
| Command ==>
|
| Inquiry Step: IOA DASD Skew Analysis
|
| Specify global data selection exit          ==> N (Y/N/R)
|
| Specify alternate input file                ==> N (Y/N/R)
| Save selected input observations            ==> N (Y/N/R)
|
|
|
|
|
|
|
|
|
|

```

Figure 2-111. IOA DASD Skew Analysis Report Extended Options Menu Panel

This Extended Options Menu panel allows you to specify that:

- o Global data selection statements are to be defined.
- o Non-CA MICS input files are to be defined.
- o Non-CA MICS output files are to be defined.

The default for all three options is N (No). The R (Reset) option allows you to return the option to its original contents. For detailed information about the extended options, press the help key (usually PF1) when you are at the Extended Options Menu.

2.4.5.4.3 Global Data Selection Panel

If you enter Y (Yes) in the Specify global data selection exit field, the Global Data Selection Panel, which is shown in Figure 2-112, is displayed. This panel allows you to enter SAS statements that will perform special data selection functions in addition to the normal DATE/HOUR/ZONE option on the main inquiry panel.

Special considerations come into play for 3380D/E, 3380J/K and 3390 devices. These devices may use either the Device Level Select (DLS) architecture or Device Level Select Enhanced (DLSE). DLS allows any two devices on a string to be active at any time and DLSE, which requires a 3990-2 or 3990-3 control unit, allows any four devices to be active. Only 3390J/K and 3390 devices can take advantage of DLSE. Both DLS and DLSE eliminate an internal path structure that can be analyzed via the substring function. Hence, the 3380D/E/J/K devices must be represented to the DASD Skew Analysis routine as being different than the earlier 3380 devices. Prior to RMF version 4.2, the data only identifies these devices as 3380, and you must tell the inquiry that you are using 3380D/E or 3380J/K devices. Furthermore, for 3380J/K attached to 3990-2 or 3390-3 control units operating in DLSE mode, the DASD skew report cannot determine this situation prior to RMF version 4.2. Consequently, it is necessary to tell the inquiry that 3990 control units are being utilized so that the inquiry reports on an assumption of DLSE architecture.

The recommended method of doing this in a "pre-RMF 4.2" environment is to specifically identify, via the global data selection exit, each 3380D/E or 3380J/K device address. For instance, if you have configured 3380Js on device addresses 480 through 48F and 580 through 58F, and they are attached to a 3990 control unit, you would create code like that shown in Figure 2-112.

```

/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| |                                           |
| Inquiry Step: IOA DASD Skew Analysis      |
| |                                           |
| Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy  |
| |                                           |
| Cmd SAS Statements for: Specifying global data selection |
| - ----- |

```

```
|  - %MACRO SELECT;                                     |
|  - IF DEVADDR = : '48' OR DEVADDR =: '58' THEN DO;    |
|  -     DEVTYPE = '3380J';_____                    |
|  -     CTLUNAME = '3990-2';_____                    |
|  - END;_____                                         |
|  - _____                                         |
|  - _____                                         |
|  - _____                                         |
|  - _____                                         |
|  - _____                                         |
|  - _____                                         |
|  - %MEND SELECT;                                     |
| ***** BOTTOM OF DATA *****                     |
| \-----/                                           |
```

Figure 2-112. SAS IOA DASD Skew Analysis Report Global Data Selector Panel

This panel allows you to specify a SAS file that is not a part of the standard CA MICS online database as input to the IOA Analysis routines. You must specify one or more indicated SAS file names and an MVS data set name. If you do not enclose the data set name in single quotes, your TSO prefix (or user ID) is added as the first portion of the data set name.

The Alternate Input File Panel requires that all named CA MICS equivalent files be specified. If a named file does not exist, you may code any name so long as there is no file in your SAS library by that name (e.g., NONE). In this case, HARPCA and SCPLCA are only present in an MVS/XA environment. If you are not running MVS/XA, then you will have to enter these two files as NONE.

This panel allows you to specify a SAS file that is not a part of the CA MICS online database as output from the IOA Analysis routines. Use this option if you want to save a copy of the data that went into the report. You must enter one or more SAS file names for the equivalent CA MICS file(s) and the MVS data set name that will contain the SAS file(s). You can read the same data into a later inquiry by specifying the SAS file and data set names on the Specify alternate input file extended option.

The Save Selected Input Observations panel requires that all named CA MICS equivalent files be specified. In this case, HARPCA and SCPLCA are only present in an MVS/XA environment. If you are not running MVS/XA, then you must enter these two files as NONE.

2.4.6 DASD Multiple Volume Analysis Inquiry

Many installations are migrating to double density devices to conserve floor space and reduce the cost per byte of their secondary storage space. However, you must give some thought to the selection of the volumes to be mapped onto the dual density devices, since significant performance problems can result if busy volumes are mapped together.

The following sections discuss this problem.

- 1 - Standard Output
- 2 - Technique Tutorial
- 3 - Case Study
- 4 - Component Operation

2.4.6.1 Standard Output

The DASD Multiple Volume Analysis Report allows you to select the volumes that have the minimum possible contention to avoid performance problems. Figures 2-115 and 2-116 show typical inquiry panels for the multiple volume analysis report and plots. Groups of up to four devices can be mapped together through the specification of 2, 3, or 4-way mapping as a parameter on the inquiry panel. The default method is the pairing of devices.

Be very careful to ensure that all volumes being analyzed have a complete set of data from each system they are shared to.

```

/----- IOA DASD Multiple Volume Analysis Report -----\
| Command ==>                                         |
| Composing CA MICS Inquiry: IOA34 - DASD Migration Study |
| Time-span ==> DAYS (DETAIL/DAYS)                   Cycle(s) ==> 01 - 03 |
| Data base ==> B (PRIMARY) _                         _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
| Num of VOLs to be combined ==> 2 (2/3/4) Plot contention index ==> Y (Y/N) |
| Volume table method: ** Specified ** : Least active method: |
|   Analyze or Monitor ==> A (A/M) : Device type ==> _____ (3350/3380) |
|   Analyze unlike VOLs ==> N (Y/N) : Num of VOLs ==> ____ (1-100) |
|   Specify volume table ==> Y (Y/N/R) : |
| Data Selection:                                     Allow run-time override |
|   Date ==> _____ - _____ - _____ ==> N (Y/N) |
|   Hour ==> 08 - 17 _ - _ _ _ _ _ _ ==> N (Y/N) |
|   Zone ==> _ - _ - _ - _ - _ _ _ _ _ _ ==> N (Y/N) |
|   Sysid ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> N (Y/N) |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-115. IOA DASD Multiple Volume Analysis Report Inquiry Panel

A description of the inquiry fields is presented in Section 2.4.6.4.1 in the discussion on component operation.

If you code Y (Yes) in the Specify volume table field, the Volume Serial Number Table panel is displayed. This panel allows you to code up to 100 volumes, keeping in mind that the number of volumes should match with the 2, 3, or 4 way mapping you have requested (e.g., you specified if 3-way mapping, then the number of volumes should be divisible by 3).

```

/----- Volume Serial Number Table -----\
| Command ==>                               |
| Inquiry Step: I0A Multiple DASD Volume Analysis |
|                                         |
|           Volume serial numbers (Max 100 in any order) |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| PACK01  PACK02  PACK03  _____  XXXACC  XXXACT  _____  _____  _____  _____ |
| BSDIS1  BSDIS3  BSDIS4  _____  BSDS01  BSDS02  BSDS05  BSDS06  _____  _____ |
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
\-----/

```

Figure 2-116. Volume Serial Number Table Panel

The DASD Multiple Volume Analysis Report for the inquiry field entries that are shown above is shown in Figure 2-117. The following items are shown in the figure:

INTERVAL START:	The time and date of the first HARDVA File observation included in the analysis.
INTERVAL END:	The time and date of the last HARDVA File observation included in the analysis.
DURATION:	The total number of hours summarized by report, rounded up to the nearest hour.
VOLSER1:	The volume serial number of the first disk volume to be mapped onto the device.
DEVICE TYPE1:	The device type of the first disk volume to be mapped onto the device.
VOLSER2:	The volume serial number of the second disk volume to be mapped onto the device.
DEVICE TYPE2:	The device type of the second disk volume to be mapped onto the device.
VOLSER3:	The volume serial number of the third disk volume to be mapped onto the device. This will be displayed only when 3-way or 4-way mapping is specified.

DEVICE TYPE3:	The device type of the third disk volume to be mapped onto the device. This will be displayed only when 3-way or 4-way mapping is specified.
VOLSER4:	The volume serial number of the fourth disk volume to be mapped onto the device. This will be displayed only when 4-way mapping is specified.
DEVICE TYPE4:	The device type of the fourth disk volume to be mapped onto the device. This will be displayed only when 4-way mapping is specified.
AVERAGE PRODUCT:	The average of the product of the device busy values (DVAPCACT) for the volumes for all of the hours processed. The average product value is defined as the contention index in the technique tutorial (Section 2.4.6.2). This value is the sum of values for all SYSIDs included in the analysis.
MAXIMUM SUM:	The maximum sum of the device busy values (DVAPCACT) for the volumes for any of the hours processed. This value is the sum of all SYSIDs included in the analysis.

We recommend that no volumes be combined that have a maximum sum value greater than 30 percent.

The software can make no assumptions concerning technological improvements that may tend to decrease device busy values as may be the case when the device types being mapped are of a different type than those volumes you are migrating to.

If you specify the optional plot reports, a plot of the sum of the device busy values for each of the hours will be produced for any volume serial grouping that has an average product of greater than 5 or a maximum sum of greater than 25. A sample plot report is shown in Figure 2-118. The purpose of these plots is to allow you to verify that a low average product value is not the result of a group of volume serial numbers that are zero, except for one or two hours that have a high contention index.

The plots present the sum of the device busy values of the recommended groupings of volume serial numbers (VOLSER grouping) for each hour of every day specified on the inquiry panel. Two values of this sum are plotted:

- o The sum of the device busy values when all device busy values were greater than zero. This variable (data element SUM1PCBS) is plotted using alphabetical characters.
- o The sum of the device busy values. This variable (data element SUM2PCBS) is plotted using asterisks. In the event that all of the groups were non-zero and less than 100, these characters will be hidden by the alphabetical characters. If the values exceeded 100, SUMPC1BS will be set to 100 and SUMPC2BS will be set to 110.

CA MICS Performance Manager
DASD MULTIPLE VOLUME ANALYSIS

INTERVAL START: 27DECyy:08:02
INTERVAL END : 27DECyy:17:59
DURATION : 09:00:00
SYSIDS : SYSA SYSB SYSC

VOLSER1	DEVICE TYPE1	VOLSER2	DEVICE TYPE2	AVERAGE PRODUCT	MAXIMUM SUM
-----	-----	-----	-----	-----	-----
BSDS02	3350	BSDS05	3350	0.01	1.33
BSDIS1	3350	BSDS01	3350	0.03	2.42
BSDIS4	3350	BSDS06	3350	0.05	4.58
PACK03	3350	XXXACT	3350	2.52	9.48
PACK02	3350	XXXACC	3350	65.90	50.31

Figure 2-117. IOA DASD Multiple Volume Analysis Report




Figure 2-118. Plot of DASD Contention Index

2.4.6.2 Technique Tutorial

One technique to reduce the cost and floor space requirements of DASD devices is the implementation of double density devices. One implementation of this technology is in devices such as the STC 8650 that allow two 3350 volumes to be mapped to a single device. Similarly, two 3350 volumes can be mapped to a single 3380 device, two single density 3380 volumes can be mapped to a single 3380E (double density 3380), or up to three 3380D devices can be mapped to a single 3380K, or up three 3380K devices can be mapped to a single 3390, and so on.

As the name implies, these devices can contain two, or three times, as much data as a standard device of the same type. Migration to 3380 devices capable of holding more than one volume's data involves the physical movement of data residing on the devices to be mapped together to the 3380 device. The implementation of double density 3350 type devices such as the STC 8650 was more complex. To allow the operating system to use the extra storage space, two apparently unique volume serial numbers are mapped to the double density device by specialized utilities provided by the vendor. If STC 8650s are in use, the operating system believes that it is using two unique physical devices, while in reality it is using a single physical device with only one disk arm. In either implementation, if the volumes that are mapped together have high I/O activity during the same time period(s), severe arm contention (i.e., delays) result.

Therefore, you must carefully select the volumes to be mapped together to ensure that there is not a high probability of the devices being active during the same time periods. The algorithm presented in this section allows you to select volume pairs that have the lowest probability of contention. If groupings of more than two (3-way or 4-way) volumes are requested, the algorithm is expanded to reflect this.

To introduce this algorithm, we present the following problem. Consider the average percent busy values (data element DVAPCACT) for the six 3380 volumes listed below. Note that if multiple SYSIDs shared the volumes, the value of DVAPCACT would be the sum of the values for each SYSID. These observations represent the activity of the six volumes for a single hour as reported in the DAYS timespan of the HARDVA File.

VOLSER	DVAPCACT
-----	-----
DATA01	3.0
DATA02	4.0
DATA03	1.0
DATA04	7.0
DATA05	2.0
DATA06	5.0

The first step in the process is the calculation of the contention and summation indices. The contention index is defined as the product of the device busy percents for all of the unique volume serial number pairs. Stated as an equation, the contention index, C , is calculated as:

$$C(i,j) = DVAPCACT(i) * DVAPCACT(j), \text{ for } j > i \text{ (Equation 10)}$$

The summation index is defined as the sum of the device busy percents for all of the unique volume serial number pairs. Stated as an equation, the summation index, S , is calculated as:

$$S(i,j) = DVAPCACT(i) + DVAPCACT(j), \text{ for } j > i \text{ (Equation 11)}$$

$$S(i,j) = 0 \text{ if } DVAPCACT(i) = 0 \text{ or } DVAPCACT(j) = 0$$

Note that the summation index is zero unless both volumes are active during the hour. Also, most of the elements of the matrix need not be computed, because the values for (i,j) are always identical to the values for (j,i) .

Figure 2-119 shows the values of the contention and summation indices for the six volume serial numbers.

	D	D	D	D	D	D
	A	A	A	A	A	A
	T	T	T	T	T	T
C (i,j)	A	A	A	A	A	A
	0	0	0	0	0	0
	1	2	3	4	5	6
DATA01	.	12	3	21	6	15
DATA02	.	.	4	28	8	20
DATA03	.	.	.	7	2	5
DATA04	.	.	.	.	14	35
DATA05	.	.	.	.	.	10
DATA06	.	.	.	.	.	.
S (i,j)						
DATA01	.	7	4	8	5	8
DATA02	.	.	5	11	6	9
DATA03	.	.	.	8	3	6
DATA04	.	.	.	.	9	12
DATA05	.	.	.	.	.	7
DATA06	.	.	.	.	.	.

Figure 2-119. Contention and Summation Index Values

Based on the contention index values, the best pair of volume serial numbers to map together would be DATA03 and DATA05, and the worst pair would be DATA04 and DATA06. However, this choice is based on data collected for a single hour. During some other time periods, there might be severe contention between the volumes.

To evaluate the potential contention between the volumes for multiple hours, we simply expand the definition of the contention index to:

```

C (i,j,k)
| | |
| | |
| | +----- is the hour number, k. It is incremented
| |           for every unique CA MICS DAY HOUR pair.
| |
| |
| |
| +----- is the pointer to the second volume serial
|           number of the pair.
|
+----- is the pointer to the first volume serial
           number of the pair.

```

For example, if 20 volumes were studied for 12 hours per day over a ten-day period, i would range from 1 to 19, j from 2 to 20, and k from 1 to 120.

After we have calculated the values of the contention and summation indices for all potential pairs for all hours, we can calculate two additional values for each potential pair. They are the average contention index and the maximum summation index. They are defined as:

$$\bar{C}(i,j) = \frac{\sum_{k=1}^K C(i,j,k)}{K} \quad (\text{Equation 12})$$

$$\text{MAXS}(i,j) = \text{MAX}(S(i,j,k)) \text{ for all } k \quad (\text{Equation 13})$$

where K is the total number of hours studied.

We can then determine volume pairing by selecting the pairs that have the lowest value of the average contention index. Since an average can easily hide one or two large values in a large sample (i.e., a study of many hours), we can use the maximum summation index as a check value. The inquiry plots the summation indices for all of the hours studied for any proposed volume pairs with an average contention index of greater than 5 or a maximum summation index greater than 25.

For 3350 type volumes, plug-compatible manufacturers have offered two methods of mapping the volumes onto the double density 3350 type drive. They are:

- o Back-to-Back - Maps the 555 cylinders of each 3350 volume into a contiguous block on the dual density device. This minimizes seek time as long as most seeks are within a single volume.
- o Interleaved - Interleaves the 555 cylinders of each 3350 volume onto the dual density device; i.e., cylinder 1 volume 1, cylinder 1 volume 2, cylinder 2 volume 1, cylinder 2 volume 2, etc. This mapping technique is recommended for volumes that are expected to have very light concurrent activity.

We recommend the following selection criteria for choosing the mapping technique:

- o Back-to-Back - For any pair whose average contention index is less than 50 and maximum summation index is less than 30.
- o Interleaved - For any pair whose average contention index is greater than 50 and less than 100 and maximum summation index is less than 50.

Do not map any pair whose contention index is greater than 100 or that has a significant number of hourly summation indices that are greater than 30 percent (as shown in the plots).

2.4.6.3 Case Study

For this case study, we used the DASD Multiple Volume Analysis Report to analyze 16 volume serial numbers for mapping on eight dual density devices. The volume serial numbers were PACK01, PACK02, PACK03, XXXACC, XXXACT, BSDIS1, BSDIS3, BSDIS4, BSDL03, BSDS01, BSDS02, BSDS05, BSDS06, BSDL05, FERC01, and BSDIS4. The time period was from 7:00 a.m. to 6:00 p.m. each day for the past three days. The CA MICS database containing the installation's hardware data was denoted by database identifier B. The inquiry fields for the study are shown in Figures 2-120 and 2-121.

```

/----- IOA DASD Multiple Volume Analysis Report -----\
| Command ==>                                           |
| Composing CA MICS Inquiry: IOA39 - Volume Grouping Analysis Case Study |
| Time-span ==> DAYS (DETAIL/DAYS)                      Cycle(s) ==> 01 - 03 |
| Data base ==> B (PRIMARY)                               |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
| Num of VOLs to be combined ==> 2 (2/3/4) Plot contention index ==> Y (Y/N) |
| Volume table method: ** Specified ** : Least active method: |
| Analyze or Monitor ==> A (A/M) : Device type ==> _____ (3350/3380) |
| Analyze unlike VOLs ==> N (Y/N) : Num of VOLs ==> ____ (1-100) |
| Specify volume table ==> Y (Y/N/R) : |
| Data Selection:                                         Allow run-time override |
| Date ==> _____ - _____ - _____ ==> N (Y/N) |
| Hour ==> 07 - 18   - - - - - ==> N (Y/N) |
| Zone ==> - - - - - ==> N (Y/N) |
| Sysid ==> _____ ==> N (Y/N) |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-120. Case Study - DASD Multiple Volume Analysis Report Inquiry Panel

The DASD Multiple Volume Analysis Report produced by the inquiry panels shown above is provided in Figure 2-122. The corresponding DASD Contention Index Plot is provided in Figure 2-123. The volume serial numbers of the suggested pairing are presented in the plot report subtitle. The volumes are PACK02 and XXXACC.

```

                                CA MICS Performance Manager
                                DASD MULTIPLE VOLUME ANALYSIS

INTERVAL START: 27DEC88:07:44
INTERVAL END   : 04JAN89:17:59
DURATION       :      31:00:00
SYSIDS        : SYSA SYSB SYSC


```

VOLSER1	DEVICE TYPE1	VOLSER2	DEVICE TYPE2	AVERAGE PRODUCT	MAXIMUM SUM
-----	-----	-----	-----	-----	-----
BSDIS3	3380	BSDL05	3380	0.01	1.86
BSDS02	3380	BSDS05	3380	0.01	1.33
BSDIS1	3380	BSDS01	3380	0.03	2.42
BSDIS4	3380	BSDS06	3380	0.05	4.58
PACK01	3380	BSDL03	3380	0.06	2.33
PACK03	3380	FERC01	3380	2.52	9.48
PACK02	3380	XXXACC	3380	65.90	50.31

Figure 2-122. Case Study - IOA DASD Multiple Volume Analysis Report




Figure 2-123. Case Study - Plot of DASD Contention Index

2.4.6.4 Component Operation

The inquiry allows you to define the CA MICS database cycles to use as input, the time periods to be analyzed, and the reports to be produced. Figures 2-124 through 2-132 show the inquiry fields used to produce the DASD Multiple Volume Analysis Report.

The following sections describe each control panel:

- 1 - DASD Multiple Volume Analysis Report Inquiry Panel
- 2 - Extended Options Menu Panel
- 3 - Global Data Selection
- 4 - Alternate Input File Panel
- 5 - Save Selected Input Observations Panel
- 6 - Volume Serial Number Table Panel
- 7 - The 2-Way Volume Serial Number Table Panel
- 8 - The 3-Way Volume Serial Number Table Panel
- 9 - The 4-Way Volume Serial Number Table Panel

2.4.6.4.1 DASD Multiple Volume Analysis Report Inquiry Panel

The DASD Multiple Volume Analysis Report is produced by executing the DASD Multiple Volume Analysis Report Inquiry panel (see Figure 2-124). This report normally analyzes and recommends DASD volumes to be combined together on a higher density device. You can produce an alternative report that monitors existing double density volumes (8650s) or proposed groupings (2, 3, or 4-way) through specifying the MONITOR option and specifying which volumes are to be grouped.

The following is a discussion of the fields on this panel.

Timespan:	The specification of whether the data is to be taken from the DETAIL or DAYS timespan.
Cycle(s):	One or more cycles from the CA MICS files that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a cycle from the CA MICS file usually contains a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided by the DATE and HOUR control parameters.
Database:	The single-character CA MICS identifiers for the databases you want to access.
Allow Runtime Override:	The indication of whether the database identifier and/or the Cycles specification can be overridden at inquiry execution time.
Number of VOLs to be combined:	The specification that determines whether groups containing two, three, or four devices must be combined and reported.
Plot contention index:	The specification of whether the plot of the inter-volume contention index will be produced for each recommended grouping.
Volume Table Analysis Method:	The following three fields combine to specify which volumes are to be analyzed or monitored.
Analyze or Monitor:	ANALYZE specifies that a recommendation report is to be produced. MONITOR specifies that the groupings you specified will be used in the

analysis.

Analyze unlike VOLS:	Specifies that groups containing unlike device types (such as a 3380 with a 3350) must be allowed. This is not recommended and will occur only if explicitly specified.
Specify Volume Table:	Specifies that a table of VOLSERs is to be defined.
Least Active Volume Analysis Method:	The following two fields instruct the inquiry to analyze the specified number of volumes having the least activity.
Device type:	The device type (3350, 3380) of the DASD that is being combined to a higher density device.
NUM of VOLS:	The number of DASD volumes to be analyzed having the least activity.
Date:	The day(s) used for the analysis. This field is normally used with the IOA Analysis inquiry to drop the extra day from a cycle due to the timing of the SMF dump.
Hour:	An hour or range of hours for which data is reported.
Zone:	The specification that zone selection is being applied. If zone 1 definition is 8:00 a.m. to 6:00 p.m., daily reports covering that time period can be specified by using the ZONE field set to 1.
Specify extended options:	The specification that allows you to supply additional report options. These options include special SAS data selection statements, non-standard input, and output files. These options are discussed in the sections that follow.

The operation of the time selection control parameters depends on definitions made during the installation of CA MICS. If you are not sure about the zone definitions for your installation or the CA MICS database update

2.4.6.4.2 Extended Options Menu Panel

If you enter Y (Yes) in the Specify extended options field, the Extended Options Menu, shown in Figure 2-125, is displayed.

```
/----- Extended Options Menu -----\
| Command ==>
|
| Inquiry Step: IOA Multiple DASD Volume Analysis
|
| Specify global data selection exit          ==> N (Y/N/R)
|
| Specify alternate input file                ==> N (Y/N/R)
| Save selected input observations            ==> N (Y/N/R)
|
|
|
|
|
|
|
|
```

Figure 2-125. DASD Multiple Volume Analysis Report Extended Options Menu

This Extended Options Menu panel allows you to specify that:

- o Global data selection statements are to be defined.
- o Non-CA MICS input files are to be defined.
- o Non-CA MICS output files are to be defined.

The default for all three options is N (No). The R (Reset) option allows you to return the option to its original value. For detailed information about the extended options, press the help key (usually PF1) when you are at the Extended Options Menu.

2.4.6.4.3 Global Data Selection Panel

If you enter Y (Yes) in the Specify global data selection exit field, the Global Data Selection Panel, which is shown in Figure 2-126, is displayed. This panel allows you to enter SAS statements that will perform special data selection functions in addition to the normal DATE/HOUR/ZONE option on the main inquiry panel.

As we discussed in the technique tutorial (see Section 2.4.6.2), the I/O Analysis inquiry merges the RMF data from multiple systems by assuming that all of an installation's volume serial numbers are unique. In the event that this assumption is invalid for your installation, you will have to specify a global data selection exit.

To illustrate the use of this exit, consider an installation that is comprised of two processors, SYSA and SYSB. Although most of the volume serial numbers in the installation are unique, the system resident volumes of each of the systems are named SYSRES. To prevent the inquiry from assuming that the two unique SYSRES packs are only one shared volume, you would code the exit as is shown below.

```

/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step: IOA Multiple DASD Volume Analysis          |
| Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy   |
| Cmd  SAS Statements for: Specifying global data selection |
| - ----- |
| - %MACRO SELECT; |
| - IF SYSID='SYSA' THEN DO; |
| -   IF VOLSER='SYSRES' THEN VOLSER='SYSAAA'; |
| - END; |
| - IF SYSID='SYSB' THEN DO; |
| -   IF VOLSER='SYSRES' THEN VOLSER='SYSBBB'; |
| - END; |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - %MEND SELECT; |
| ***** BOTTOM OF DATA ***** |
\-----/

```

Figure 2-126. DASD Multiple Volume Analysis Report Global Data Selection Panel

2.4.6.4.4 Alternate Input File Panel

If you enter Y (Yes) in the Specify alternate input file field in the Extended Options Menu, the Alternate Input File panel, shown in Figure 2-127, is displayed.

[illegible]

Figure 2-127. DASD Multiple Volume Analysis Report Alternate Input File Panel

This panel allows you to specify a SAS file that is not a part of the standard CA MICS online database as input to the IOA Analysis routines. You must specify one or more indicated SAS file names and an MVS data set name. If you do not enclose the data set name in single quotes, your TSO prefix (or user ID) is added as the first portion of the data set name.

2.4.6.4.6 Volume Serial Number Table Panels

Use the Volume Serial Number Table panels to specify the DASD volumes to be analyzed by the DASD Multiple Volume Analysis Inquiry.

Any number of volumes can be specified, as long as the total number of volume serial numbers is: divisible by the number of volumes to be mapped together as specified on the DASD Multiple Volume Analysis Report inquiry panel (i.e., if a 3-way report request is specified, the number of volumes specified should be divisible by three).

If different device types (3330, 3350, and 3380) are input, the quantity of each device type's volume serial numbers should be divisible by the report request type, because recommended groupings will normally be of a like device type.

[illegible]

Figure 7-129. IOA DASD Multiple Volume Analysis Volume Serial Number Table Panel

If the DASD Multiple Volume Analysis inquiry is to be used to MONITOR existing combined volumes or to measure volume groupings of your choice, then you specify the volume serial numbers that are to be monitored on the 2, 3, or 4-WAY Volume

2.4.6.4.9 The 4-WAY Volume Serial Number Table Panel

Use the 4-WAY Volume Serial Number Table panel to define quadruples of volumes that are already combined or are to be combined. The DASD Multiple Volume Analysis inquiry will MONITOR these volumes and report how well they are matched. Entry fields are grouped into fours for your convenience.

```

/----- 4-WAY Volume Serial Number Table -----\
| Command ==>                                     |
| Inquiry Step: IOA Multiple DASD Volume Analysis |
|                                                  |
|              Volume serial numbers in groups of 4 |
|              -----                          ----- |
| 1= ----- 5= ----- |
| 9= ----- 13= ----- |
| 17= ----- 21= ----- |
| 25= ----- 29= ----- |
| 33= ----- 37= ----- |
| 41= ----- 45= ----- |
| 49= ----- 53= ----- |
| 57= ----- 61= ----- |
| 65= ----- 69= ----- |
| 73= ----- 77= ----- |
| 81= ----- 85= ----- |
| 89= ----- 93= ----- |
| 97= ----- |
|                                                  |
\-----/

```

Figure 2-132. IOA DASD Multiple Volume Analysis 4-WAY Volume Serial Number Table Panel

2.4.7 Tape Drive Simulation Analysis Inquiry

The unavailability of sufficient tape drives can adversely affect performance, especially during periods of heavy batch job activity. If sufficient tape drives are not available, batch work is delayed, and operators may have to vary tape drives from one system to another.

The Tape Drive Simulation Analysis inquiry allows you to analyze time periods that have significant tape drive unavailability problems. You can also model the effects of moving existing work to different time periods, adding additional tape usage, and tape drive improvements. Thus, you can investigate both the effects on performance of rescheduling existing jobs and the capacity implications of added tape drive usage.

The following sections describe the facility's operation and reports:

- 1 - Standard Output
- 2 - Technique Tutorial
- 3 - Case Study
- 4 - Component Operation

2.4.7.1 Standard Output

Tape Drive Simulation Analysis provides both summary and detail level graphs and reports displaying periods of allocation. Examples of each are included in this section.

The summary level emphasizes time periods having tape drive allocations which exceed exception limits set by the user and displays only work showing tape drive usage. The detail level shows all workloads independent of tape drive use during the period being analyzed, so that you can make judgments about which time periods might be candidates for more tape drive usage, or which work might be delayed or accelerated in order to better balance the usage of tape drives.

The following sections discuss:

- 1 - Report Titles
- 2 - Summary Level Reports
- 3 - Detail Level Reports

2.4.7.1.1 Report Titles

The title of each report indicates the analysis features you have selected. These and other analysis features which you can select are more fully explained in Section 2.4.7.4, Component Operations. The title for each of the example reports in this section indicates that the following features were used:

STEP HIGH

WATERMARK/ Either the JES3 step high water mark allocation
STEP method or the JES2 step allocation method was
ALLOCATION: analyzed.

SNAPSHOT

METHOD USED/

PRORATING Either the snapshot algorithm was used or the
METHOD USED: one which uses prorating.

WORK ADDED: The user built a model including additional
 hypothetical workloads which allocated tape
 drives.

WORK MOVED: The user built a model by moving existing
 workloads to different time periods.

MACROS

ACTIVE: Indication that you have supplied macros into
 the control stream in order to modify the
 standard report or data.

2.4.7.1.2 Summary Level Reports

Summary level reports have two parts:

- 1 - Summary Chart
- 2 - Summary Exception Reports

2.4.7.1.3 Detail Level Reports

Detail level reports also have two parts, the Detail Workload Chart and Detail Workload Report.

DETAIL WORKLOAD CHART

The Detail Workload Chart shown in Figure 2-135 is similar to the Summary Chart, but distinguishes between allocation of tape drives by workload type. By default, workload type is defined as BATCH, STARTED TASKS, APPC, and TSO for unmodified observations. If an observation has been modified (moved), then the workload type is defined as M. If the observation has been added, then its workload type is identified as A.

The following table describes what type of workloads the symbols used in the Detail Workload Chart represent:

SYMBOL USED	WORKLOAD REPRESENTED
A	Workloads hypothetically added by the analyst.
B	Batch work currently allocating tapes at the time being reported.
M	Actual workloads hypothetically moved to a different time interval by the analyst.
S	Started tasks currently allocating tapes at the time being reported.
T	TSO work dynamically allocating tapes at the time being reported.
P	APPC work dynamically allocating tapes at the time being reported.

```

CA MICS Performance Manager

TAPE DRIVE ANALYSIS

DETAIL WORKLOAD CHART
TAPE ALLOCATION BY TYPES OF WORK RUNNING DURING THE ANALYSIS PERIOD
STEP HIGH WATERMARK MODEL - PRORATING METHOD USED - WORK ADDED - WORK MOVED - SELECTION MACRO ACTIVE

DATE ANALYZED=27APRyy HOUR OF DAY=1

BAR CHART OF SUMS

TESTTM  TIME ANALYZED
1:00:00  |
1:05:00  |AAAAAAAAAAAAAAAAAAAAAM
1:10:00  |AAAAAAAAAAAAAAAAAAAAAMMM
1:15:00  |AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAMMM
1:20:00  |AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABMMMM
1:25:00  |AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABBBBBBM
1:30:00  |AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABBBBBB
1:35:00  |AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABBBBB
1:40:00  |AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABBBB
1:45:00  |AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABBBBBB
1:50:00  |AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABBBBB
1:55:00  |AAAAAAAAAAAAAAAAAAAAABB
-----+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
      2  4  6  8 10 12 14 16 18 20 22 24 26 28 30 32 34 36

PGMMXNTA SUM

SYMBOL SYM          SYMBOL SYM          SYMBOL SYM
A  ADDED  WORKLOAD   B  BATCH WORKLOAD   M  MOVED  WORKLOAD

```

Figure 2-135. Tape Drive Simulation Analysis - Detail Workload Chart
DETAIL WORKLOAD REPORT

The second portion of the detail level reports, the Detail Workload Report, is a listing of all workloads during the period being analyzed. This report appears below in Figure 2-136. It is nearly identical in overall appearance to the Summary Exception Report described above; however, it is not limited to workloads using tape drives and will be much more voluminous.

The Detail Workload Report does contain one additional item of interest not found in the Summary Exception Report.

WRKLOAD: BATCH, STC, TSO, APPC, ADDED, or MOVED
 to indicate the source of this particular
 workload.

Refer to the description of the Summary Exception Report above for more information about the format and content of these two reports.

CA MICS Performance Manager							8
TAPE DRIVE ANALYSIS							
DETAIL WORKLOAD REPORT							
ALL WORKLOADS RUNNING DURING THE ANALYSIS PERIOD							
STEP ALLOCATION MODEL - SNAPSHOT METHOD USED -WORK ADDED -							
-----DATE ANALYZED=07SEPyy TIME ANALYZED=7:12:00-----							
JOB ID	READER TIME STAMP	STEP NBR	PROGRAM NAME	START TIME STAMP			
PSG01703	07SEPyy:15:55:31.58	4	UNJABTAB	07SEPyy:06:56:33.64			
SYS053DA	07SEPyy:07:03:51.93	1	SASLPA	07SEPyy:07:04:07.84			
TEP100N1	07SEPyy:07:09:59.65	1	WCSGENER	07SEPyy:07:10:10.25			

TESTTM							
JOB ID	END TIME STAMP	SYS ID	STEP ALLOC. TIME	WRKLOAD	TAPE DEVICES ALLOC.	PROGRAMS EXECUTING	
PSG01703	07SEPyy:07:12:37.85	XA01	0:00:05.53	BATCH WORKLOAD	2.00	1	
SYS053DA	07SEPyy:07:32:50.12	XA01	0:00:09.78	BATCH WORKLOAD	1.00	1	
TEP100N1	07SEPyy:07:15:28.53	XA04	0:00:01.23	BATCH WORKLOAD	1.00	1	

TESTTM					4.00	3	
-----DATE ANALYZED=07SEPyy TIME ANALYZED=7:14:00-----							
JOB ID	READER TIME STAMP	STEP NBR	PROGRAM NAME	START TIME STAMP			
SYS053DA	07SEPyy:07:03:51.93	1	SASLPA	07SEPyy:07:04:07.84			
TEP100N1	07SEPyy:07:09:59.65	1	WCSGENER	07SEPyy:07:10:10.25			

TESTTM							
JOB ID	END TIME STAMP	SYS ID	STEP ALLOC. TIME	WRKLOAD	TAPE DEVICES ALLOC.	PROGRAMS EXECUTING	
SYS053DA	07SEPyy:07:32:50.12	XA01	0:00:09.78	BATCH WORKLOAD	1.00	1	
TEP100N1	07SEPyy:07:15:28.53	XA04	0:00:01.23	BATCH WORKLOAD	1.00	1	

TESTTM					2.00	2	

Figure 2-136. Tape Drive Analysis - Detail Workload Report

2.4.7.1.2.1 Summary Chart

The Summary Chart is shown in Figure 2-133. One of these charts is produced for each period (hour or date) analyzed. The date and hour being reported in each chart are printed in the subtitle near the top of the chart.

The vertical axis of the chart shows the time intervals during which analyses were done. The horizontal axis represents the number of tape drives allocated.

The vertical reference line represents a user-defined exception limit. Any interval during which the number of tape drives allocated exceeds this limit is more fully reported in the Summary Exception Report, described below.

The following table describes what type of workload the symbols used in the Summary Chart represent:

SYMBOL USED	WORKLOAD REPRESENTED
A	Workloads hypothetically added by the analyst.
C	Actual work currently running at the time being reported.
M	Actual workloads hypothetically moved to a different time interval by the analyst.

Chapter 2: PERFORMANCE REPORTING ANALYSIS 363

2.4.7.1.2.2 Summary Exception Reports

The second portion of the summary level reports consists of the Summary Exception Report, which is shown in Figure 2-134. This report is produced only if two conditions are met:

- o The user defines an exception limit.
- o One or more time periods reported exceed this limit.

An exception limit of 0 causes all observations to print.

This report contains one or more subreports, one for each time interval during which the exception limit was exceeded. Each subreport is preceded by its own subheading containing the following identifiers:

DATE

ANALYZED: The date of this exception period.

TIME

ANALYZED: The time of this exception period.

STANDARD ITEMS

The report contains the following standard items:

JOB ID: The name of the job being analyzed. Any hypothetical jobs added during the analysis will have their JOB ID begin with an asterisk (*).

READER
TIME STAMP: The date and time that the job was read into the system.

STEP NBR: The number of the job step being analyzed.

PROGRAM
NAME: The name of the program being executed by this job step.

START

TIME STAMP: The time that the job step began executing.

END

TIME STAMP: The time that the job step ended its execution.

SYS ID: The SYSID of the system in which the job step executed.

STEP ALLOC. The amount of time the step waited for device

TIME:	allocation by MVS to complete. This time represents periods when the step was eligible for execution, but required I/O devices (that is, DASD, tape, etc.) were not yet allocated for its use. This condition can affect the Tape Drive Simulation Analysis inquiry adversely. The Tape Drive Simulation Analysis inquiry assumes a device to be allocated from step start (data element STARTTS) to step end (data element ENDTS).
ADDED OR MOVED:	The indication of whether the job step represents actual work currently running at the time being reported (CURRENT), a workload hypothetically added (ADDED), or actual work that has been hypothetically moved to a different time interval (MOVED).
TAPE DEVICES ALLOC.:	The number of tape drives allocated by this job step. This always represents the number of tape drives actually needed by the operating system for this job step as represented by the data element PGMMXNTA.

The standard items displayed are controlled by a panel containing lists of data element names. Any data element from the CA MICS BATPGM type file is eligible for display. In addition, the data elements JOBINQTM and JOBFETTM from the BATJOB File are available. Also, any user-computed data elements are eligible for display.

OPTIONAL ELEMENTS

In addition to the standard items, the following optional items are printed, depending on the inquiry fields specified.

If you choose JES Version 3 (main device scheduling) the number of tape drives reserved by the allocation method is calculated and reported.

HIGH WATER ALLOC.:	The number of tape drives that would be required by this job step if the high water mark option of JES3 Main Device Scheduling was enabled. This value is computed in the Tape Drive Simulation Analysis inquiry in the data element HIGHWATR.

If the Analysis Method field specifies PRORATE and you choose JES Version 2, then the number of tape drives used within an

interval is calculated and displayed.

TAPE DRIVES PRORATED (MVS): The number of tape drives used by this job step in the interval. This number is the percentage of tape drives prorated to the interval. For example, let us assume an interval of 60 minutes and a step that executed for 45 minutes within the interval. If the step had a PGMMXNTA value of 4, the resulting number of tape drives prorated to the interval/step combination would be 3 ($4 \times (45/60)$). The element is computed in the Tape Drive Simulation Analysis inquiry in the data element PROTAPE.

If the Analysis Method field specifies PRORATE and you choose JES Version 3, then, in addition to the above data elements, the number of drives prorated to the interval for the step is computed, substituting the number of drives reserved under JES3 highwater marks allocation for the number allocated under MVS. A description follows:

TAPE DRIVES PRORATED (MDS): The number of tape drives allocated or reserved by the step within the interval. This number is a percentage computed by the Tape Drive Simulation Analysis inquiry in the data element PROHIGH. It uses the methodology documented above, under TAPE DRIVES PRORATED (MVS), substituting the number of drives reserved or allocated as documented under the item HIGH WATER ALLOC for the number of drives allocated (PGMMXNTA).

Use of the USER Modeling exit causes additional data elements to be displayed. This exit provides a method of adjusting step execution times based on user-supplied values. The use of this facility causes the following items to be displayed:

ORIGINAL END TIMESTAMP: The unmodified end timestamp. If any adjustment to execution time is made in a standard fashion, this element (ORGENDTS) is displayed.

If mount time adjustment is made, then, in addition to the unmodified end timestamp, the following items are displayed:

REEL DEVICES ALLOCATED: The number of REEL tape devices allocated. This element is computed as the total number of tape devices (PGMMXNTA) minus the number of CTRG devices (PGMN3480).

CTRG DEVICES The number of cartridge devices allocated
 ALLOCATED: (PGMN3480).

JOB

ADJUSTMENT The running total of step adjustment time that
 TIME: has taken place in prior steps of the job.

SPECIFIED The specified mount time value in seconds.
 AVERAGE This item is supplied to the Tape Drive
 MOUNT TIME: Simulation Analysis inquiry by the USER Modeling
 exit and causes the step's end timestamp to be
 adjusted.

TAPE The sum of specific and non-specific tape mounts
 MOUNTS: as found in the BATPGM File data elements
 PGMSTMT and PGMNSTMT. These data elements are
 available at the MVS/SE2 levels and above.

COMPUTED Tape mounts multiplied by the specified average
 MOUNT TIME: mount time. This value is added to the end
 timestamp. If the specified average mount time
 was negative, the computed mount time is
 negative as well.

If any tape connect time adjustment is made using the
 documented features provided in the USER Modeling exit, then
 the original connect time is displayed in addition to the
 other elements.

ORIGINAL The time the step tape drives were connected for
 TAPE CONNECT the purpose of data transfer (PGMCTAPE). This
 TIME: element is only available in MVS/XA operating
 systems.

Connect time adjustment is done by a percent increase or
 decrease, or by specification of an average I/O time. The
 percent method can only be used if connect time is available
 (MVS/XA element). Connect time and EXCPs are computed for
 REEL devices as the total (PGMCTAPE and PGMETAPE) minus the
 observed values of the CTRG counters (PGMC3480 and PGME3480).
 Adjustment of connect time by device type can then proceed.
 The following items are thus method (percent vs. average I/O
 time) and device (REEL vs. CTRG) dependent.

ORIGINAL These elements contain the unmodified connect
 CONNECT TIME time for each type of tape device. When the
 REEL DEVICES/ average I/O time method is used, the elements
 ORIGINAL are computed. When the percent method is

CONNECT TIME used, the REEL time is computed as the total
CTRG DEVICES: tape device connect time minus the CTRG
connect time.

SPECIFIED REEL These items represent the percentage
CONNECT % adjustment factor supplied to the tape
ADJUSTMENT/ Drive Simulation Analysis inquiry via the
SPECIFIED CTRG USER Modeling exit for each device type.
CONNECT % Adjustment factors are specified by the
ADJUSTMENT: device type (REEL or CTRG) independently.
This decimal factor is then multiplied by
connect time to derive the adjustment amount.

SPECIFIED These items represent the average I/O time to
AVERAGE REEL complete one EXCP as supplied to the Tape
I/O TIME/ Drive Analysis Simulation inquiry via the
SPECIFIED USER Modeling exit. These factors are
AVERAGE CTRG specified by device type (REEL or CTRG). The
I/O TIME: EXCPs for the device type are multiplied by
the factor to derive an adjustment amount.

REEL EXCPs/ These elements are displayed when the average
CTRG EXCPs: I/O time method is used. EXCPs for
REEL devices are computed as the total tape
EXCPs (PGMETAPE) minus the CTRG EXCPs
(PGME3480). When the average I/O time method
is used, EXCPs by device type are multiplied
by the supplied times.

ADJUSTED
CONNECT TIME-
REEL DEVICES/ These elements are displayed when connect time
ADJUSTED adjustment has occurred for the device times,
CONNECT TIME- and contain the amount of time to adjust the
CTRG DEVICES: end timestamp of a step.

CA MICS Performance Manager									
TAPE DRIVE ANALYSIS									
SUMMARY EXCEPTION REPORT									
TAPE WORKLOADS RUNNING DURING TAPE-CONSTRAINED PERIODS									
STEP ALLOCATION MODEL - SNAPSHOT METHOD USED - WORK ADDED									
-----DATE ANALYZED=07SEPyy TIME ANALYZED=8:00:00-----									
JOB ID	READER TIME STAMP	STEP NBR	PROGRAM NAME	START TIME STAMP	END TIME STAMP	SYS ID	STEP ALLOC. TIME	ADDED OR MOVED	TAPE DEVICES ALLOC.
*W	0	0		07SEPyy:06:04:00.00	07SEPyy:08:12:00.00		0	ADDED	3.00
PSG049AA	07SEPyy:07:57:50.07	4	IEBGENER	07SEPyy:07:58:27.72	07SEPyy:08:02:46.50	XA01	0:00:00.94	CURRENT	1.00
SYS053BA	07SEPyy:07:39:34.84	1	SASLPA	07SEPyy:07:39:49.36	07SEPyy:08:12:11.18	XA01	0:00:13.62	CURRENT	1.00
-----									-----
TESTTM									5.00
-----DATE ANALYZED=07SEPyy TIME ANALYZED=8:02:00-----									
JOB ID	READER TIME STAMP	STEP NBR	PROGRAM NAME	START TIME STAMP	END TIME STAMP	SYS ID	STEP ALLOC. TIME	ADDED OR MOVED	TAPE DEVICES ALLOC.
*W	0	0		07SEPyy:06:04:00.00	07SEPyy:08:12:00.00		0	ADDED	3.00
MCC001DD	07SEPyy:08:00:56.92	4	SASLPA	07SEPyy:08:01:38.03	07SEPyy:08:19:37.46	XA01	0:00:05.10	CURRENT	1.00
PSG049AA	07SEPyy:07:57:50.07	4	IEBGENER	07SEPyy:07:58:27.72	07SEPyy:08:02:46.50	XA01	0:00:00.94	CURRENT	1.00
SYS053BA	07SEPyy:07:39:34.84	1	SASLPA	07SEPyy:07:39:49.36	07SEPyy:08:12:11.18	XA01	0:00:13.62	CURRENT	1.00
-----									-----
TESTTM									6.00
-----DATE ANALYZED=07SEPyy TIME ANALYZED=8:04:00-----									
JOB ID	READER TIME STAMP	STEP NBR	PROGRAM NAME	START TIME STAMP	END TIME STAMP	SYS ID	STEP ALLOC. TIME	ADDED OR MOVED	TAPE DEVICES ALLOC.
*W	0	0		07SEPyy:06:04:00.00	07SEPyy:08:12:00.00		0	ADDED	3.00
MCC001DD	07SEPyy:08:00:56.92	4	SASLPA	07SEPyy:08:01:38.03	07SEPyy:08:19:37.46	XA01	0:00:05.10	CURRENT	1.00
PSG00690	07SEPyy:08:01:43.57	4	TMSGW	07SEPyy:08:02:04.68	07SEPyy:08:04:46.23	XA01	0:00:02.43	CURRENT	1.00
SYS053BA	07SEPyy:07:39:34.84	1	SASLPA	07SEPyy:07:39:49.36	07SEPyy:08:12:11.18	XA01	0:00:13.62	CURRENT	1.00
-----									-----
TESTTM									6.00

Figure 2-134. Tape Drive Simulation Analysis - Summary Exception Report

2.4.7.2 Technique Tutorial

Tape drives, unlike DASD, are usually not shared between systems. Generally, tape devices are physically attached (cabled) to multiple systems, but are physically switched to only one at a time.

Physical paths to tape devices are established by using switches (toggles) on the control units. Both 3420 and 3480 control units have channel adapter switches. The 3420 control units also have separate toggle switches for each device.

It is becoming more common to "share" the tape subsystem by having some of the tape devices physically switched to two systems but logically offline to one or the other. This allows all systems access to the tape subsystem, while preventing concurrent access to any single device. Software access to the drives is accomplished by the use of the VARY ONLINE operator command.

The 3480 tape subsystem can function in one of two modes:

- o Full Function
- o 3420 Compatibility

In Full Function mode, the drives are assigned a processor by the use of the VARY ONLINE command. The command causes values to be placed in the pathing array within the control unit.

After assignment to a system, no other system can access the drive unless the VARY command employs the SHR option. The SHR option allows the drive to be assigned to multiple systems. In this shared mode, tape devices must be controlled by software (such as JES3 or other products) to prevent concurrent access.

In 3420 Compatibility mode, pathing arrays do not exist. The only way to prevent concurrent access is through the use of the channel adapter switches.

As you can see from the above description, performance in the area of batch throughput can be adversely affected not only by the total number of tape drives available, but also by the number of drives online to each system in a multiprocessor configuration. Tape drive capacity and performance prediction is an area often overlooked due to the complexity and changeable nature of workloads and configurations in a

multiprocessor environment.

A common operational safeguard is to place the majority of the tape devices on the primary tape processing system and ensure that tape job class initiators are enabled on that system. This prevents the physical switching of tape devices from one system to another when a job or started task on another system requests more tape devices than are available to it. For this reason, you must take special care to ensure that the true maximum tape drive limit is known for each system prior to running the Tape Drive Simulation Analysis inquiry. Generally, the inquiry should be run separately for each processor in a multiprocessor complex.

As mentioned above, JES3 provides the capability to manage tape devices which are online to more than one system, unlike the standard MVS allocation scheme, which is only aware of device status on one system. The JES3 facility that provides device management is the Main Device Scheduler (MDS). MDS attempts to satisfy the resource requirements of a job in the context of a globally controlled environment, so that job execution may proceed without allocation delays.

Standard MVS allocation is done at the step level as step execution begins. MDS, on the other hand, considers the requirements for each step of a job, within the context of other currently scheduled work, prior to the start of job execution. MDS reserves devices, thus making them unavailable for other work, then allocates them at the step level as in standard MVS.

RMF/SMF data sources do not currently report JES3 reservations, although the effect is similar to MVS allocation in that the device is not available while it is reserved. For this reason, MDS device requirements are higher than MVS device requirements. On the positive side, MDS minimizes tape device contention between jobs.

Under MVS allocation, job steps are competing for resources, with the losers (those steps requesting unavailable resources) either being cancelled by the operations staff or waiting for the resource(s) to become available.

This waiting period is costly in that the waiting job is holding other valuable resources such as an initiator, an address space, data sets, and possibly devices. These held resources can affect overall system utilization adversely. MDS attempts to avoid this situation by not allowing jobs to be passed to MVS for execution until their allocation

requirements can be met or "set up."

JES3 "setup" occurs while the job is in the JES3 address space (the only resource used during setup is JES3 queuing space). MDS reserves devices, then requests and verifies the mounting of initial volumes on each device before the job can be selected for execution. If "defer" is specified, volume mount requests are not issued at setup.

The type of setup that MDS employs can be specified on the `//*MAIN` JES3 control statement included in the job's JCL stream, or via the JES3 STANDARDS initialization statement. Using the initialization method is the common approach.

The types of setup that MDS can employ are:

- o Job
- o High Water Mark
- o Explicit

Job or High Water Mark setup can be specified at JES3 initialization. Explicit is very similar to Job setup, but can only be specified via the `//*MAIN` control statement. These types of setup are discussed in detail below.

JOB SETUP

Job setup attempts to premount all volumes for all the steps in a job prior to passing the job to MVS for initialization. For example, a five-step job that requires three tape data sets in each step causes MDS to reserve 15 tape devices, ensure that these devices are mounted, and then pass the job to MVS for initialization. Job setup is thus quite expensive in terms of the drives required.

HIGH WATER MARK SETUP

High Water Mark setup is more efficient in device allocation. Under this method, the greatest number of devices that any step requires is determined prior to job execution, and that number of devices is reserved for the job. In the above example, three drives are reserved and the volumes required by the first step premounted. Under both Job and High Water Mark setup, drives no longer required are released as the step ends.

A more complex example illustrates High Water Mark setup and the releasing of drives. Job X has six steps, where step 1 requires one tape drive, step 2 requires five, step 3

requires no drives, step 4 requires four, step 5 requires one, and step 6 requires two. This information is displayed in the following chart:

	STEP #					
	1	2	3	4	5	6
Drive Requirements	1	5	0	4	1	2
High Water Mark Reservations	5	5	4	4	2	2
Job Setup Reservations	13	12	7	7	3	2

As shown above, High Water Mark setup ensures that steps will not wait for device allocation after execution begins, and is much more efficient in its use of tape devices. The Drive Requirements line of the chart can be equated to standard MVS allocation if the drives are available when the step begins execution.

EXPLICIT SETUP

A third method, known as Explicit setup, can be specified via the `//*MAIN JES3` control statement. Only the premounting of drives is affected under Explicit setup. Device reservation defaults to the method used in Job setup.

The type of setup to employ or whether to employ MDS setup is dependent on your installation's needs and requirements. For more information, consult the following IBM manuals: JES3 System Programming Library: Installation, Planning, and Tuning; JES3 Introduction; and JES Overview.

The Tape Drive Simulation Analysis inquiry allows you to examine the utilization of your tape drive subsystems during periods that have significant tape drive availability problems. It also permits you to develop simple models to analyze the effects on the tape subsystem of moving or adding workloads. By calculating the allocation of tape drives repeatedly over sufficiently short intervals, this inquiry can aid you in ensuring efficient tape drive utilization when scheduling existing or new work.

You can use Tape Drive Analysis inquiry for the following purposes:

- o Provide graphical and tabular reports on the status of

the tape subsystem.

- o Model the effects of change on the subsystem.
- o Aid in capacity planning and performance tuning of the tape subsystem.

The Tape Drive Simulation Analysis inquiry provides two algorithms for reporting tape drive allocation levels. One provides a precise snapshot of drive allocation levels at each measured time interval. The other prorates drive usage according to the percentage of the interval during which any particular job step was executing.

The first method may be most helpful when you want to obtain a moment-by-moment report of tape drive allocation patterns. The second method is particularly helpful in answering capacity-related questions.

The following sections discuss the operation of the Tape Drive Analysis inquiry in terms of:

- 1 - Data Limitations
- 2 - Computational Methods Available
- 3 - Analyzing JES2 and JES3 Drive Allocation
- 4 - Exception Reporting
- 5 - Modeling
- 6 - User Modeling Exit

2.4.7.2.1 Data Limitations

The Tape Drive Simulation Analysis inquiry is based on the SMF data that is event driven (that is, the data record is produced when the event occurs, such as step or interval ending). The Tape Drive Analysis inquiry assumes that all events occurring during a report interval are input and selected for the report.

This assumption may not be valid, unless the report interval time being modeled includes the data from the report specification time to a time period when all work (steps) running during the report interval has ended and been processed by the CA MICS MVS Batch and Operations Analyzer. For example, if a step spans two generations of SMF data and only the first generation data is known to the Tape Drive Simulation Analysis inquiry, then the report spanning known data may not include some of the data for that step.

This limitation is common to SMF event-driven data and can adversely affect computations based on timespans. To limit the extent of this problem and improve event reporting, interval records (Type 30 SMF records) can be enabled and input to CA MICS. When interval accounting is enabled in both SMF and CA MICS, this problem is limited to the loss of information for one interval at most, rather than for the entire job step, as could happen without Type 30 SMF records.

The existence of allocation time in a step record represents a problem in the modeling process and for the assumption that tape drives are allocated from step start to step end. Allocation time as found in the CA MICS data element PGMALCTM cannot be broken down to delineate which device allocation contributed time to it (that is, what the allocation time was caused by).

Tape drives may be freed by job control prior to step end. This invalidates the assumption that drives are allocated from step start to step end and can be corrected by coding a user exit.

This type of limitation can best be explained by the following example. Let us assume a configuration of four drives and a workload of two steps (from two separate jobs), where Step 1 used three drives and Step 2 subsequently requested two drives. In this scenario, Step 2 first goes in allocation wait. The records processed by the drive inquiry could reflect a total of five drives for the period when both Step 1 and Step 2 were executing concurrently.

If, on the other hand, Step 1 had freed one drive prior to Step 2's request, and the data was input without modification, the drive inquiry would again reflect five for the periods when Step 1 and Step 2 were concurrently executing.

2.4.7.2.2 Computational Methods Available

The inquiry obtains data from the DETAIL levels of the BATPGM File (for batch jobs) and from the BAT_ST File (for started tasks). Optionally, TSO work from the optional BAT_TS file or APPC information from the BAT_TP file can be included if relevant to the study. Since the number of tape drives allocated or reserved/allocated to a particular job can change from one job step or step interval to another, and since the amount of time during which each particular step has exclusive use of tape drives must be known for this type of analysis, the inquiry requires this detailed level of data.

Using this data as input, the inquiry then performs a "stepped" computation of the number of drives allocated or reserved/allocated during each time interval analyzed. The number of drives allocated during the interval represents the number of drives allocated from the step start to step end and under MVS control. The number of drives reserved/allocated represents the number allocated (maximum of JES3 MDS high water mark reservations (HIGHWATR) or the MVS allocated number (PGMMXNTA)) from step start to step end.

You can use either one of two methods to perform this analysis:

- o Snapshot Method
- o Prorating Method

SNAPSHOT METHOD

Using the snapshot method, if, for example, the interval clock value of the inquiry is set to two minutes, then the inquiry determines the number of allocated tape drives at each two-minute interval from the beginning to the end of the analysis period specified. In this way, it calculates snapshots of the allocation status of the tape drive subsystem.

This method has the advantage of providing an exact sample of tape drive activity at the time of each sample interval. Its disadvantage is that it misses any job step which starts after one sample interval and completes prior to the next sample interval.

To minimize this potential effect of overlooking jobs steps shorter than the sampling interval, the inquiry does not

permit this method to use a sampling interval longer than two minutes. If you do not specify a sample interval, the default time interval used is one minute. Analyses of a variety of data have shown that one minute sample intervals are usually satisfactory for this method of analysis.

If you wish, you can specify sampling intervals shorter than one minute. Although shorter sampling intervals improve the accuracy of this method slightly, they also lead to greater use of resources.

Reports and charts produced under the snapshot method represent intervals within an hour.

PRORATING METHOD

The prorating method does not overlook tape drive utilization by even very short job steps. It also credits each job step which begins or ends part way through a sampling interval with partial use of any tape drives that it has allocated during this interval. This method provides an accurate calculation of the average number of drives allocated during any size sampling interval.

It is important to remember, when using this method, that excessively long intervals tend to mask the real-world impact of short-term peaks in tape drive utilization. This effect is characteristic of any analysis based upon averages, and is not unique to the inquiry.

You need to determine what the maximum useful sampling interval is, based upon your understanding of your installation's workload mix and operational objectives. Analyzing the same data using several different sampling intervals may help you determine what is the largest sampling interval which still retains sufficient detail about peaks and valleys in the allocation pattern for tape drives.

Reports and charts produced under this method are displayed by day if the specified interval (INTV) is greater than 60 minutes; otherwise they are displayed by hour within day.

2.4.7.2.3 Analyzing JES2 and JES3 Drive Allocation

The inquiry can accommodate both JES2 installations and JES3 installations (using the JES3 Main Device Scheduler (MDS) High Watermark facility). JES2 systems allocate and free tape drives on a job step by job step basis.

JES3 systems can use this same process or, at the beginning of the job, they can reserve the maximum number of drives that will be needed by any job step, and then reduce this number of reserved drives after the step requiring this maximum number finishes. Under this process, a new somewhat smaller number of drives, representing the maximum number of drives required in any following job step, are reserved. This process continues until the last step of the job finishes. This high water mark device scheduling ordinarily results in somewhat greater tape drive requirements than does the step by step process.

To aid detail analyses and assist in rescheduling work, the DETAIL level BATJOB File data elements JOBCLASS, JOBMCLS, JOBINQTM, and JOBFETTM are available for display.

2.4.7.2.4 Exception Reporting

The summary reports produced by the Tape Drive Simulation Analysis inquiry will be most useful to you if you specify the chart reference line option. This is a high tape allocation limit beyond which exception reporting will occur. This could be considered to be the practical limit for long-term drive allocation.

The value given to this parameter would ordinarily be somewhat smaller than the number of physical tape drives actually present in the installation. Generally, the number should represent the number of tape drives online to the system. RMF data contained in the CA MICS HARDVA Files can be analyzed to determine this number. Setting the chart reference line to 0 causes all data to be printed in the summary reports.

2.4.7.2.5 Modeling

The report requests can be accompanied by several other types of optional exits that expand the straightforward historical analyses into models. These models have tape drive allocations that have been altered in one of several ways from that actually reported in the CA MICS database.

The reports issued under this modeling feature are identical to the standard reports, except that each report is clearly identified as a MODEL report, instead of as an ANALYSIS report.

The standard modeling features permit two types of modifications of tape drive allocation:

- o Specifying additional tape drive use
- o Moving existing tape use to a different time period

SPECIFYING ADDITIONAL TAPE DRIVE USE

You can specify additional hypothetical use of tape drives via the Add workload observations extended option. This facilitates study of the effects of additional tape drive use beyond that which already occurs.

MOVING EXISTING TAPE USE TO A DIFFERENT TIME PERIOD

The second feature, moving existing tape use to a different time period, allows you to specify a new reader timestamp, RDRTS, for one or more jobs via the Move workload observations extended option. The RDRTS represents the date and time when the job first entered the system. If one or more jobs are moved, the new STARTTS and ENDTS are calculated as follows for each job:

$$\text{ENDTS}(\text{new}) = \text{ENDTS}(\text{old}) + (\text{RDRTS}(\text{new}) - \text{RDRTS}(\text{old}))$$

$$\text{STARTTS}(\text{new}) = \text{STARTTS}(\text{old}) + (\text{RDRTS}(\text{new}) - \text{RDRTS}(\text{old}))$$

You can use this feature, for example, to model the shift of a "tape hog" job to a different time slot.

MODELING CONCERNS

The modeling technique used is very simple. As a linear unconstrained model, it does not take into consideration any resource scarcity or queuing. Such a model can still be very

helpful when attempting to schedule batch workloads or plan for hardware requirements. There are several implications of using a linear unconstrained model:

- o Queuing delays due to I/O path contention are not included in the model.
- o There is no initiator or memory queuing. If a job is moved or added, it is assumed by the model that it has all other resources necessary to run it, including an available initiator and adequate memory. In addition, if one job is moved, the job that actually followed it in an initiator would not have its starting time automatically altered.

These implications are in general no different from those of any modeling; you should understand the model's limits and be sure that the model represents your real world with sufficient accuracy to be helpful.

More complex modeling, including predicting the effects of shifting to 3480 drives, can be accomplished using a general purpose exit that is described in Section 2.4.7.2.6.

2.4.7.2.6 USER Modeling Exit

The USER modeling exit occurs after the data is extracted from the DETAIL CA MICS files. It gives you the ability to make additional changes to the original data for more complex modeling requirements.

This exit is especially useful when modeling the effects of jobs using standard technology tape drives versus 3480 drives or a mixture of the two technologies.

At this exit point, you have access to and may modify any of the variables extracted from the BATPGM, BAT_ST, BAT_TS, BAT_TP, BATJOB, and BAT_JS Files. However, because the Tape Drive Simulation Analysis inquiry performs all its calculations using only the STARTTS, ENDTS, and PGMMXNTA (the Step Start Time Stamp, the Step End Time Stamp, and Magnetic Devices Allocated), any modifications must ultimately be reflected in these three variables or there will be no effect on the Tape Drive Simulation Analysis inquiry.

You can use this exit by selecting the Specify USER modeling exit on the Extended Options Menu panel (see Section 2.4.7.4 for a detailed description of the operation of a Tape Drive Simulation Analysis inquiry).

```

/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step: I/OA Tape Drive Simulation Analysis |
| Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy |
| Cmd SAS Statements for: Specifying USER exit |
| - ----- |
| _ - %MACRO PERUSEL; |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| _ - %MEND PERUSEL; |
| ***** BOTTOM OF DATA ***** |
\-----/

```

Figure 2-137. User Modeling Exit Panel

When you enter Y (Yes) to the Specify USER modeling exit on the Extended Options Menu, the User Modeling Exit Panel, shown in Figure 2-137, is displayed. You can use this exit to modify a job step's start and end timestamps based on tape drive connect time or mount time. Connect time can be computed based on a supplied EXCP time. Connect time adjustments can be selectively applied to a step by device type (3480 or 3420). Mount time adjustment cannot be done by device type. When mount time adjustment is made on a job that uses both 3480 and 3420 drives, a warning message is issued.

You should regard results from the use of this exit with caution. The inquiry makes no attempt to adjust other factors based on the effects of changes specified in this exit.

2.4.7.3 Case Study

For this case study, we used the Tape Drive Simulation Analysis inquiry to determine peaks of daily tape drive usage, determine the effects of tape drive over-allocation on two critical jobs, and model the implementation of the movement of one of these jobs to a period of time when tape drives were available. Based on the results of the analysis, we planned an additional study detailing steps that allocated tape drives but did not transfer any tape data.

The site in question used a shared pool of 32 3420 type devices. Two drives were 800/1600 BPI (bytes per inch) devices and the rest were 1600/6250 BPI. The vast majority of the tape work used tapes at 6250 BPI. A midweek IPL was usually scheduled Wednesday at midnight. Scheduled tape drive maintenance generally occurred during prime shift when tape drive and batch usage was low. JES2 was in effect with a purchased job scheduling system. Prime shift hours 8-18 were primarily devoted to CICS, TSO, and batch test work. Normally CICS was brought down at 20:00, at which time production batch work was heavily scheduled.

The study was triggered by the failure to meet objectives in production work schedules for critical tape jobs. The initial reports were run over a period of several weeks to determine if tape drive over-allocation was the cause of delays, when the peaks of allocation fell, and if a consistent daily/hourly pattern was apparent. Figure 2-138 shows the control parameters we used for the last of these daily reports.


```

/----- IOA Tape Drive Simulation Analysis -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: IOA57 - Tape Drive Analysis Case Study |
| Time-span:    DETAIL                               Cycle(s) ==> 01 - 01 |
| Data base ==> A (PRIMARY) _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Include batch pgms ==> Y (Y/N)      Include system tasks ==> N (Y/N) |
| Include TSO users  ==> N (Y/N)      Merge SMF JOB info.  ==> N (Y/N) |
| Include APPC TPs   ==> N (Y/N)                                     |
| Report date ==> 13NOVyy (ddmonyy)      Report hours  ==> 0 - 23 (hh) |
| Report level ==> SUMMARY (SUMMARY/DETAIL) JES version ==> 2 (2/3) |
| Anal. method ==> PRORATE (SNAPSHOT/PRORATE) Sampling intv. ==> 60 (min) |
| Max. tapes  ==> 30 (Blank or nn)                                     |
| Chart reference line and summary report limit for printing ==> 0 (000-999) |
| Data Selection :                                     |
|   Hour Selection ==> __ - __ (hh-hh) |
|   Sysid Selection ==> _____ |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-138. Case Study - Tape Drive Simulation Analysis Inquiry Panel (Report 1)

The Tape Analysis Summary Chart and an extract of the report for hour 23 is provided in Figure 2-146. The pattern displayed in the chart is typical of those days when failures to meet scheduling deadlines were reported. Relatively high step allocation times in hour 23 were typical. Two of the jobs that were listed as critical production work, CNSPD032 and UCC7LOGS, displayed extremely high allocation delay times and relatively high (2) drive demands. System and job output for these two jobs confirmed that the allocation delays were the result of tape drive demands. Hours 3-8 consistently showed low tape drive usage periods, as did prime shift and other periods when CICS was active.

The summary reports also showed a large number of job steps that were extremely short in duration allocating tape drives. This type of work displayed zero drives prorated in the summary report due to the length of the interval (1 hour or 3600 seconds) and the very short durations of the steps (less than 18 seconds). (The report display length for the prorated number of drives is limited to two decimal positions, thus these values were displayed as zero.) It appeared that these job steps either conditionally required tape drives, or were using tape when DASD would have been more efficient.

Investigation of the JCL of a few of these jobs revealed that in many cases the JCL was structured to ensure at job start that allocation delays would not be experienced by subsequent steps requiring a drive. In other words, unneeded tape drive allocations were implemented by previous steps. In effect, tape drive reservation was being implemented manually through the JCL. We concluded that a separate study of these steps must be conducted and that, in the future, allocations of this type must be regarded as a standards violation in production JCL.

Obviously, the movement of work from the tape-constrained period shown in the peak of hours 22 and 23 to a less utilized period would cause the allocation delays experienced by work in hour 23 to decrease and shorten the durations of both CNSPD032 and UCC7LOGS. The scheduling requirements for both jobs were analyzed. The only requirement for CNSPD032 was completion by 8:00; UCC7LOGS, however, could not be moved.

Hours 3 through 6 appeared to be a good candidate for the movement of CNSPD032. Figure 2-139 shows the entries we made in the inquiry fields to analyze these hours in detail to

determine the best starting point for CNSPD032. An extract of the charts showing hour 3 is shown in Figure 2-147. The reports for this time frame are not shown. The charts and reports for hours 4-6 showed results very similar to hour 3 after 3:24:00, in that no time period exceeded 10 drives.

```

/----- IOA Tape Drive Simulation Analysis -----\
| Command ==>                                     |
|                                                    |
| Composing CA MICS Inquiry: IOA58 - Tape Drive Analysis Case Study |
|                                                    |
| Time-span:    DETAIL                               Cycle(s) ==> 01 - 01 |
| Data base ==> A (PRIMARY) _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
|                                                    |
| Include batch pgms ==> Y (Y/N)      Include system tasks ==> N (Y/N) |
| Include TSO users  ==> N (Y/N)      Merge SMF JOB info.  ==> N (Y/N) |
| Include APPC TPs   ==> N (Y/N)                                     |
|                                                    |
| Report date ==> 13NOVyy (ddmonyy)      Report hours ==> 3 - 6 (hh) |
| Report level ==> DETAIL (SUMMARY/DETAIL) JES version ==> 2 (2/3) |
| Anal. method ==> SNAPSHOT (SNAPSHOT/PRORATE) Sampling intv. ==> 2 (min) |
| Max. tapes ==> 30 (Blank or nn)                                     |
| Chart reference line and summary report limit for printing ==> 26 (000-999) |
|                                                    |
| Data Selection :                                     |
|   Hour Selection ==> 03 - 06 (hh-hh)                                     |
|   Sysid Selection ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
|                                                    |
| Specify extended options ==> N (Y/N)                                     |
|                                                    |
\-----/

```

Figure 2-139. Case Study - Tape Drive Simulation Analysis Inquiry Panel (Report 2)

Based on this output, we decided to run two more inquiries to model the movement of CNSPD032 to 3:24:00 and to model effects of reduced allocation times for both jobs. The inquiry fields for each run are shown in Figures 2-140 through 2-193 and 2-194 through 2-196.

```

/----- IOA Tape Drive Simulation Analysis -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: IOA59 - Tape Drive Analysis Case Study |
| Time-span:    DETAIL                               Cycle(s) ==> 01 - 01 |
| Data base ==> A (PRIMARY) _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Include batch pgms ==> Y (Y/N)      Include system tasks ==> N (Y/N) |
| Include TSO users  ==> N (Y/N)      Merge SMF JOB info.  ==> N (Y/N) |
| Include APPC TPs   ==> N (Y/N)                                     |
| Report date  ==> 13NOVyy (ddmmmyy)      Report hours   ==> 22 - 23 (hh) |
| Report level ==> SUMMARY (SUMMARY/DETAIL) JES version  ==> 2          (2/3) |
| Anal. method ==> SNAPSHOT (SNAPSHOT/PRORATE) Sampling intv. ==> 2      (min) |
| Max. tapes   ==> 30      (Blank or nn)                                     |
| Chart reference line and summary report limit for printing ==> 26 (000-999) |
| Data Selection :                                     |
|   Hour Selection ==> 22 - 23 (hh-hh)                                     |
|   Sysid Selection ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Specify extended options ==> Y (Y/N)                                     |
\-----/

```

Figure 2-140. Case Study - Tape Drive Simulation Analysis Inquiry Panel (Report 3)

Because we coded Y (Yes) in the Specify extended options field, the Extended Options Menu was displayed, which enabled us to chose the Specify USER modeling exit and Move workload observations options.


```

/----- IOA Tape Drive Simulation Analysis -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: IOA62 - Tape Drive Analysis Case Study |
| Time-span:    DETAIL                               Cycle(s) ==> 01 - 01 |
| Data base ==> A (PRIMARY) _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Include batch pgms ==> Y (Y/N)      Include system tasks ==> N (Y/N) |
| Include TSO users  ==> N (Y/N)      Merge SMF JOB info.  ==> N (Y/N) |
| Report date  ==> 13NOVyy (ddmonyy)      Report hours   ==> 03 - 06 (hh) |
| Report level ==> SUMMARY (SUMMARY/DETAIL) JES version  ==> 2      (2/3) |
| Anal. method ==> SNAPSHOT (SNAPSHOT/PRORATE) Sampling intv. ==> 1      (min) |
| Max. tapes   ==> 30      (Blank or nn) |
| Chart reference line and summary report limit for printing ==> 0      (000-999) |
| Data Selection : |
|   Hour Selection ==> 03 - 06 (hh-hh) |
|   Sysid Selection ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Specify extended options ==> Y (Y/N) |
\-----/

```

Figure 2-143. Case Study - Tape Drive Simulation Analysis Inquiry Panel (Report 4)

```

/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step: I0A Tape Drive Simulation Analysis |
| Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy |
| Cmd SAS Statements for: Specifying USER exit |
| - ----- |
| - %MACRO PERUSEL; |
| - IF JOB = 'UCC7LOGS' OR JOB = 'CNSPD032' THEN DO; |
| -   ENDTS=ENDTS-PGMALCTM; |
| - END; |
| - %MEND PERUSEL; |
| ***** BOTTOM OF DATA ***** |
| |
| |
| |
\-----/

```

Figure 2-144. Case Study - User Modeling Exit (Report 4)

```

/----- Move Workload Observations -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step: I0A Tape Drive Simulation Analysis |
| Report date: 13NOVyy   Report hours: 03 - 06 |
| Line Cmds: I Insert  D Delete  C Copy  M Move  R Reset |
| |
|           Old Start Time           New Start Time |
| Cmd   Jobname   (hh : mm : ss . th)   (hh : mm : ss . th) |
| - ----- |
| - CNSPD032      21 : 04 : 48 . 43      03 : 24 : 00 . 00 |
| ***** BOTTOM OF DATA ***** |
| |
| |
| |
\-----/

```

Figure 2-145. Case Study - Move Workload Observation (Report 4)

Extracts from the charts for hours 22 and 23 are shown in Figure 2-148. As the summary report was requested with a Drive Limit of 26 and the UCC7LOGS job was modified to reduce its execution by 30:00.70 (the allocation time) in the USER exit macro, the step now ends at 22:45:23.33. Thus, the Drive Limit value was not exceeded during the time periods of the job's execution, and the job is not displayed on the report.

Figure 2-149 displays extracts from the charts and reports for hour 3 showing the effect of the movement of CNSPD032.




Figure 2-146. Case Study - Report 1 - Summary Prorate Chart/Report Extract (Part 1 of 3)

CA MICS Performance Manager					
TAPE DRIVE ANALYSIS					
SUMMARY EXCEPTION REPORT					
TAPE WORKLOADS RUNNING DURING TAPE-CONSTRAINED PERIODS					
STEP ALLOCATION ANALYSIS - PRORATING METHOD USED - MACROS ACTIVE -					
-----DATE ANALYZED=13NOVyy TIME ANALYZED=23:00:00-----					
JOB ID	READER TIME STAMP	STEP NBR	PROGRAM NAME	START TIME STAMP	
ASBNKBUP	13NOVyy:23:17:49.84	1	IKCGENER	13NOVyy:23:17:56.90	
ASBNKBUP	13NOVyy:23:17:49.84	2	IKCGENER	13NOVyy:23:19:49.27	
ASBNKBUP	13NOVyy:23:17:49.84	3	IKCGENER	13NOVyy:23:19:55.31	
BAK10SDI	13NOVyy:23:46:07.61	3	BARBS102	13NOVyy:23:47:33.91	
BEJBKTP	13NOVyy:22:01:02.58	2	IDCAMS	13NOVyy:22:37:47.26	
BEJBKTP	13NOVyy:22:01:02.58	3	IDCAMS	13NOVyy:23:10:13.97	
BEJBKTP	13NOVyy:22:01:02.58	4	IDCAMS	13NOVyy:23:46:37.59	
CCYREPR1	13NOVyy:20:05:45.75	2	IDCAMS	13NOVyy:23:44:16.56	
CCYREPR2	13NOVyy:20:05:46.05	2	IDCAMS	13NOVyy:23:43:46.80	
CHMTTBKP	13NOVyy:23:16:24.14	5	IKCGENER	13NOVyy:23:50:33.09	
CICSCMF1	13NOVyy:23:14:09.79	1	IEBGENER	13NOVyy:23:14:40.15	
CICSCMF1	13NOVyy:23:14:09.79	2	IEBGENER	13NOVyy:23:49:06.01	
CICSCMF2	13NOVyy:23:14:09.89	1	IEBGENER	13NOVyy:23:14:30.19	
CICSCMF2	13NOVyy:23:14:09.89	2	IEBGENER	13NOVyy:23:55:00.15	
CICSCMF3	13NOVyy:23:14:09.98	1	IEBGENER	13NOVyy:23:14:30.18	
CICSCMF3	13NOVyy:23:14:09.98	2	IEBGENER	13NOVyy:23:42:36.29	
CICSCMF3	13NOVyy:23:14:09.98	3	IEBGENER	13NOVyy:23:55:40.17	
CICSPK1X	13NOVyy:04:04:49.38	7	DFHSIP	13NOVyy:04:05:57.85	
CNSPD032	13NOVyy:21:04:48.43	1	CNSPD030	13NOVyy:21:12:25.08	
UCC7LOGS	13NOVyy:22:33:34.55	2	SASSHIS5	13NOVyy:22:39:59.16	

TESTTM					
TESTDATE					

Figure 2-146. Case Study - Report 1 - Summary Prorate Chart/Report Extract (Part 2 of 3)

CA MICS Performance Manager							
TAPE DRIVE ANALYSIS							
SUMMARY EXCEPTION REPORT							
TAPE WORKLOADS RUNNING DURING TAPE-CONSTRAINED PERIODS							
STEP ALLOCATION ANALYSIS - PRORATING METHOD USED - MACROS ACTIVE -							
-----DATE ANALYZED=13NOVyy TIME ANALYZED=23:00:00-----							
JOB ID	END TIME STAMP	SYS ID	STEP ALLOC. TIME	ADDED OR MOVED	TAPE DEVICES ALLOC.	MVS TAPE DRIVES PRORATED	
ASBNKBUP	13NOVyy:23:19:49.20	SYSE	0:00:58.02	CURRENT	1.00	0.03	
ASBNKBUP	13NOVyy:23:19:55.25	SYSE	0:00:01.70	CURRENT	1.00	0.00	
ASBNKBUP	13NOVyy:23:19:59.02	SYSE	0:00:01.45	CURRENT	1.00	0.00	
BAK10SDI	13NOVyy:23:50:46.07	SYSE	0:00:01.83	CURRENT	1.00	0.05	
BEJBKTP	13NOVyy:23:10:13.90	SYSE	0:00:01.92	CURRENT	1.00	0.17	
BEJBKTP	13NOVyy:23:46:37.51	SYSE	0:08:41.76	CURRENT	1.00	0.61	
BEJBKTP	13NOVyy:23:49:45.12	SYSE	0:00:01.67	CURRENT	1.00	0.05	
CCYREPR1	14NOVyy:00:03:26.44	SYSC	0:00:02.28	CURRENT	1.00	0.26	
CCYREPR2	14NOVyy:00:15:59.32	SYSC	0:00:02.71	CURRENT	1.00	0.27	
CHMTTBKP	13NOVyy:23:56:40.29	SYSE	0:01:06.87	CURRENT	1.00	0.10	
CICSCMF1	13NOVyy:23:49:05.92	SYSA	0:04:22.41	CURRENT	1.00	0.57	
CICSCMF1	14NOVyy:00:06:37.69	SYSA	0:07:43.69	CURRENT	1.00	0.18	
CICSCMF2	13NOVyy:23:55:00.07	SYSE	0:04:25.18	CURRENT	1.00	0.67	
CICSCMF2	14NOVyy:00:05:40.41	SYSE	0:05:07.72	CURRENT	1.00	0.08	
CICSCMF3	13NOVyy:23:42:36.22	SYSE	0:04:25.18	CURRENT	1.00	0.47	
CICSCMF3	13NOVyy:23:55:40.10	SYSE	0:00:02.44	CURRENT	1.00	0.22	
CICSCMF3	14NOVyy:00:02:14.23	SYSE	0:03:12.07	CURRENT	1.00	0.07	
CICSPK1X	13NOVyy:23:03:18.56	SYSB	0:00:48.49	CURRENT	1.00	0.05	
CNSPD032	13NOVyy:23:15:24.03	SYSE	0:30:00.70	CURRENT	2.00	0.51	
UCC7LOGS	13NOVyy:23:34:37.95	SYSC	0:33:52.66	CURRENT	2.00	1.15	
-----					-----	-----	
TESTTM					224.00	25.04	
TESTDATE					1356.00	199.19	
					=====	=====	
					1356.00	199.19	

Figure 2-146. Case Study - Report 1 - Summary Prorate Chart/Report Extract (Part 3 of 3)








Figure 2-149. Case Study - Report 4 - Summary Snapshot Chart/Report (Part 1 of 3)



CA MICS Performance Manager TAPE DRIVE ANALYSIS SUMMARY EXCEPTION REPORT TAPE WORKLOADS RUNNING DURING TAPE-CONSTRAINED PERIODS STEP ALLOCATION MODEL - SNAPSHOT METHOD USED - WORK MOVED - MACROS ACTIVE -									
-----DATE ANALYZED=13NOVyy TIME ANALYZED=3:35:00-----									
JOB ID	READER TIME STAMP	STEP NBR	PROGRAM NAME	START TIME STAMP	END TIME STAMP	SYS ID	STEP ALLOC.	ADDED OR MOVED	TAPE DEVICES ALLOC.
CHPTKRT2	13NOVyy:03:19:57.30	4	CHPTKIT2	13NOVyy:03:32:59.82	13NOVyy:03:41:05.85	SYSE	0:00:16.01	CURRENT	2.00
CHPUHSTD	13NOVyy:03:05:44.36	1	CHPUHST2	13NOVyy:03:05:45.00	13NOVyy:03:41:25.40	SYSC	0:00:11.94	CURRENT	1.00
CUPSWED	13NOVyy:01:41:10.55	2	CHPCUPS2	13NOVyy:01:43:18.69	13NOVyy:04:11:16.03	SYSC	0:00:33.08	CURRENT	1.00
HRM1000B	13NOVyy:03:19:48.89	2	PAY030	13NOVyy:03:19:56.56	13NOVyy:03:48:08.99	SYSB	0:00:05.04	CURRENT	1.00
MICODALY	13NOVyy:00:41:02.05	2	SAS	13NOVyy:00:45:46.41	13NOVyy:04:12:40.21	SYSC	0:00:19.29	CURRENT	1.00
TRSDTCU2	13NOVyy:02:42:53.02	5	TRSUCAS2	13NOVyy:02:49:40.80	13NOVyy:03:36:41.54	SYSB	0:00:30.13	CURRENT	1.00
CNSPD032	13NOVyy:03:24:00.00	1	CNSPM030	13NOVyy:03:31:36.65	13NOVyy:05:04:34.90	SYSE	0:30:00.70	MOVED	2.00
TESTTM									9.00
-----DATE ANALYZED=13NOVyy TIME ANALYZED=3:36:00-----									
JOB ID	READER TIME STAMP	STEP NBR	PROGRAM NAME	START TIME STAMP	END TIME STAMP	SYS ID	STEP ALLOC.	ADDED OR MOVED	TAPE DEVICES ALLOC.
CHPTKRT2	13NOVyy:03:19:57.30	4	CHPTKIT2	13NOVyy:03:32:59.82	13NOVyy:03:41:05.85	SYSE	0:00:16.01	CURRENT	2.00
CHPUHSTD	13NOVyy:03:05:44.36	1	CHPUHST2	13NOVyy:03:05:45.00	13NOVyy:03:41:25.40	SYSC	0:00:11.94	CURRENT	1.00
CUPSWED	13NOVyy:01:41:10.55	2	CHPCUPS2	13NOVyy:01:43:18.69	13NOVyy:04:11:16.03	SYSC	0:00:33.08	CURRENT	1.00
HRM1000B	13NOVyy:03:19:48.89	2	PAY030	13NOVyy:03:19:56.56	13NOVyy:03:48:08.99	SYSB	0:00:05.04	CURRENT	1.00
MICODALY	13NOVyy:00:41:02.05	2	SAS	13NOVyy:00:45:46.41	13NOVyy:04:12:40.21	SYSC	0:00:19.29	CURRENT	1.00
TRSDTCU2	13NOVyy:02:42:53.02	5	TRSUCAS2	13NOVyy:02:49:40.80	13NOVyy:03:36:41.54	SYSB	0:00:30.13	CURRENT	1.00
CNSPD032	13NOVyy:03:24:00.00	1	CNSPM030	13NOVyy:03:31:36.65	13NOVyy:05:04:34.90	SYSE	0:30:00.70	MOVED	2.00
TESTTM									9.00
-----DATE ANALYZED=13NOVyy TIME ANALYZED=3:37:00-----									
JOB ID	READER TIME STAMP	STEP NBR	PROGRAM NAME	START TIME STAMP	END TIME STAMP	SYS ID	STEP ALLOC.	ADDED OR MOVED	TAPE DEVICES ALLOC.
CHPTKRT2	13NOVyy:03:19:57.30	4	CHPTKIT2	13NOVyy:03:32:59.82	13NOVyy:03:41:05.85	SYSE	0:00:16.01	CURRENT	2.00
CHPUHSTD	13NOVyy:03:05:44.36	1	CHPUHST2	13NOVyy:03:05:45.00	13NOVyy:03:41:25.40	SYSC	0:00:11.94	CURRENT	1.00
CUPSWED	13NOVyy:01:41:10.55	2	CHPCUPS2	13NOVyy:01:43:18.69	13NOVyy:04:11:16.03	SYSC	0:00:33.08	CURRENT	1.00
HRM1000B	13NOVyy:03:19:48.89	2	PAY030	13NOVyy:03:19:56.56	13NOVyy:03:48:08.99	SYSB	0:00:05.04	CURRENT	1.00
MICODALY	13NOVyy:00:41:02.05	2	SAS	13NOVyy:00:45:46.41	13NOVyy:04:12:40.21	SYSC	0:00:19.29	CURRENT	1.00
TRSDTCU2	13NOVyy:02:42:53.02	9	IEBGNER	13NOVyy:03:36:45.21	13NOVyy:03:39:00.58	SYSB	0:00:02.71	CURRENT	1.00
CNSPD032	13NOVyy:03:24:00.00	1	CNSPM030	13NOVyy:03:31:36.65	13NOVyy:05:04:34.90	SYSE	0:30:00.70	MOVED	2.00
TESTTM									9.00

Figure 2-149. Case Study - Report 4 - Summary Snapshot Chart/Report (Part 3 of 3)

2.4.7.4 Component Operation

The Tape Drive Simulation Analysis inquiry allows you to define the CA MICS database cycles to be used as input, the time periods to be analyzed, and the reports to be produced. Figures 2-150 through 2-157 show the inquiry panels used to produce the Tape Drive Simulation Analysis Report.

The following sections describe each control panel:

- 1 - Tape Drive Simulation Analysis Inquiry Panel
- 2 - Extended Options Menu Panel
- 3 - Global Data Selection Panel
- 4 - USER Modeling Exit Panel
- 5 - Add Workload Observations Panel
- 6 - Move Workload Observations Panel
- 7 - Alternate Input File Panel
- 8 - Save Selected Input Observations Panel
- 9 - Define Report Description Element Names Panel

2.4.7.4.1 Tape Drive Simulation Analysis Inquiry Panel

Figure 2-150 shows the IOA Tape Drive Simulation Analysis Inquiry Panel. The following fields are used on that panel:

Timespan:	The specification that the data is to be taken from the DETAIL CA MICS timespan.
Cycle(s):	One or more cycles of the CA MICS files that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a cycle from the CA MICS file usually contains a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided by the REPORT DATE and HOUR parameters.
Database:	The single character CA MICS identifier for the database you want to access.
Include batch pgms:	Code a Y (Yes) if the BATPGM File is required in the analysis.
Include system tasks:	Code a Y (Yes) if the BAT_ST File is required in the analysis.
Include TSO users:	Code a Y (Yes) if the BAT_TS File is required in the analysis. Normally TSO users are restricted from tape use.
Include APPC TPs:	Code a Y (Yes) if the BAT_TP File is required in the analysis.
Merge SMF JOB info:	Code a Y (Yes) if you want reporting on the variables JOBCCLASS, JOBMCLS, JOBINQTM, or JOBFETTM.
Report date, hours, Hour Selection:	Two time selection methods are provided to allow analysis on a subset of the input data and to restrict reporting to times under current analysis. The input data is first limited to a range of data to be selected by the specification of the Hour Data Selection parameter. Range specification is required due to the nature of the data source (see Section 2.4.7.2.1, Data

Limitations).

Report date and report hour fields define the data to be presented in the reports.

Report level:	Either a SUMMARY or DETAIL report.
JES version:	The specification of whether 2 (JES2 step allocation) or 3 (JES3 high watermark) device scheduling is to be analyzed. The default is 2. The differences between these two methods of scheduling tape device allocation are described in Sections 7.7 and 2.4.7.2
Anal. method:	Either the snapshot or the prorating analytical method. The default is SNAPSHOT. These two methods are described in Section 2.4.7.2.
Sampling intv.:	The number of minutes in each sampling interval for this analysis. The default is one minute. If you choose the SNAPSHOT analytical method, an interval greater than two minutes will be flagged as an error. If you choose the PRORATE method and you specify an interval greater than 59 minutes, subsequent reports will be by day rather than by hour within day.
Max. tapes:	Specifies the maximum scale for the summary bar chart report for this analysis. This could, if you wish, be set to the maximum number of drives available. The default is the actual maximum number of drives allocated during any particular hour analyzed. Thus the default changes for each hour's bar chart.
Chart reference line and limit:	The exception limit for the summary reports for this analysis. This could, if you wish, be set to the maximum number of drives that operations experience has found can be allocated without causing extra operator activity in retaining and remounting tapes. There is no default. If this parameter is not specified, or if the data analyzed does not exceed this limit, no Summary Exception Report will be produced. If specified as 0, the Summary Exception Report will display all data for workloads using tape drives.
Sysid:	Specifies the SYSIDs of the data to be analyzed.

Specify extended options: The specification that allows you to supply additional report options. These options include special SAS data selection statements and USER modeling exits, non-standard input and output files, Add and Move workload observation, and a panel to describe additional reporting elements.

The operation of the time selection control parameters depends on definitions made during the installation of CA MICS.

```

/----- IOA Tape Drive Simulation Analysis -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: IOA69 - Tape Drive Simulation Inquiry |
| Time-span:    DETAIL                               Cycle(s) ==> 01 - 02 |
| Data base ==> A (PRIMARY) _           _           _           _ |
| Include batch pgms ==> Y (Y/N)           Include system tasks ==> Y (Y/N) |
| Include TSO users  ==> N (Y/N)           Merge SMF JOB info.  ==> N (Y/N) |
| Include APPC TPs   ==> N (Y/N) |
| Report date  ==> 01DECyy (ddmonyy)           Report hours   ==> 18 - 23 (hh) |
| Report level ==> SUMMARY (SUMMARY/DETAIL)     JES version    ==> 2         (2/3) |
| Anal. method ==> SNAPSHOT (SNAPSHOT/PRORATE) Sampling intv. ==> 1         (min) |
| Max. tapes   ==> 20         (Blank or nn) |
| Chart reference line and summary report limit for printing ==> 15 (000-999) |
| Data Selection : |
|   Hour Selection ==> __ - __ (hh-hh) |
|   Sysid Selection ==> _____ |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-150. Tape Drive Simulation Analysis Inquiry Panel

```
/----- Extended Options Menu -----\
| Command ==>
|
| Inquiry Step: IOA Tape Drive Simulation Analysis
|
| Specify global data selection exit          ==> N (Y/N/R)
| Specify USER modeling exit                 ==> N (Y/N/R)
|
| Add workload observations                    ==> N (Y/N/R)
| Move workload observations                   ==> N (Y/N/R)
|
| Specify alternate input file                  ==> N (Y/N/R)
| Save selected input observations              ==> N (Y/N/R)
|
| Define report description element names      ==> N (Y/N/R)
```

This Extended Options Menu panel allows you to specify that:

- o Global data selection statements are to be defined.
- o A USER modeling exit is to be defined.
- o Workload observations are to be added or moved.
- o Non-CA MICS input files are to be defined.
- o Non-CA MICS output files are to be defined.
- o Report description element names are to be defined.

The default for all of these options is N (No). The R (Reset) option allows you to return the option to its original contents. For detailed information about the extended options, press the help key (usually PF1) when you are at the Extended Options Menu.

2.4.7.4.3 Global Data Selection Panel

If you enter Y (Yes) in the Specify global data selection exit field, the Global Data Selection Panel, which is shown in Figure 2-152, is displayed. This panel allows you to enter SAS statements that will perform special data selection functions in addition to the normal DATE/HOUR/ZONE option on the main inquiry panel.

The SELECT MACRO can be used for special selection code based on any data element in the BATPGM type files (i.e., BATPGM, BAT_TS, BAT_ST, or BAT_TP Files). Typical uses may be to limit data analysis to a specific subset of the input, such as only production work.

```

/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step:  IOA Tape Drive Simulation Analysis      |
| Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy  |
| Cmd  SAS Statements for:  Specifying global data selection |
| - ----- |
| _ - %MACRO SELECT; |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| _ - %MEND SELECT; |
| ***** BOTTOM OF DATA ***** |
|-----/

```

Figure 2-152. Tape Drive Simulation Analysis Global Data Selection Panel

2.4.7.4.4 USER Modeling Exit Panel

This is the only modeling exit that is available in the Tape Drive Simulation Analysis inquiry. You can use this exit to select a variety of modeling capabilities whose uses are described in Section 2.4.7.2. This exit provides a window into the data after it is obtained either from the original CA MICS files or from the optional alternate input file(s). This window occurs before the Tape Drive Simulation Analysis inquiry begins its analysis or modeling using this data.

```

/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step: IOA Tape Drive Simulation Analysis |
| Line Cmds: I Insert  D Delete  R Repeat  M Move   C Copy |
| Cmd SAS Statements for: Specifying USER exit |
| - ----- |
| _ - %MACRO PERUSEL; |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| _ - %MEND PERUSEL; |
| ***** BOTTOM OF DATA ***** |
\-----/

```

Figure 2-153. Tape Drive Simulation Analysis USER Modeling Exit Panel

Because the data elements of the original CA MICS PGM type files and optionally certain data elements (JOBCLASS, JOBMCLS, JOBINQTM, and JOBFETTM) from the BATJOB and BAT_TS Files are available to the exit, it can be used for a wide range of modifications of this data prior to analysis.

There are two cautions concerning the use of this exit:

- o As mentioned in Section 2.4.7.2, the analysis makes use of only the variables STARTTS, ENDTS, and PGMMXNTA, so that any modifications to the data must ultimately be reflected in these three variables.
- o Prior to the invocation of the USER modeling exit, the following data elements are initialized to zero:
 - CHGMTTM (Change Mount Time)
 - PCTCON2 (Percent Connect 3420)
 - PCTCON8 (Percent Connect 3480)
 - AVIOTM2 (Average I/O Time, 3420)
 - AVIOTM8 (Average I/O Time, 3480)

The setting of any of these data elements in the exit to a non-zero quantity may affect the end timestamp of a step.

Following is a description of the possible uses for each of these data elements, examples, and the subsequent processing that occurs for a job. It is important to note that only the end timestamp of individual steps is affected in any of these processes. Any adjustments of a job's subsequent steps start and end timestamps must be implemented by inserting the appropriate code in the USER modeling exit.

The USER modeling exit takes place after step records from the PGM type files (BATPGM, BAT_TS, BAT_TP, and BAT_ST) and JOB type files (BATJOB and BAT_JS) data have been merged in order to supply the class data element, and after type 30 records have been summarized to the step level. At the point of invocation, the records are in the following "by" sequence: JOB, SYSID, and RDRTS (reader timestamp).

Just prior to the merge, the step records are sorted by JOB, RDRTS, and STEPNUM. Thus observations passing through the exit are in the following order: JOB, SYSID, RDRTS, and STEPNUM. Prior to invocation, the Tape Drive Simulation Analysis inquiry does the following:

ADJUSTTM represents a total job adjustment time factor that can be used to adjust subsequent steps within the same job. ADJUSTTM is the sum of all execution time adjustments made to a job. The other data elements shown above are step adjustment factors that are used to compute execution time differences. They are discussed below. ADJUSTTM can be used to alter the start and end timestamps of a step based on adjustments to earlier steps before further step adjustment is made.

Prior to the invocation of the USER modeling exit, the Tape Drive Simulation Analysis inquiry derives device type (3420 and 3480) totals for tape EXCPs, drives, and connect time by the subtraction of PGME3480, PGMN3480, and PGMC3480 from the total tape counters PGMETAPE, PGMXNTA, and PGMCTAPE, thereby creating PGME3420, PGMN3420, and PGMC3420. The device-specific data elements (those with the suffix 3420 or 3480) are set to zero if missing values are found.

Next, the inquiry checks for non-zero quantities in PCTCON2, PCTCON8, AVIOTM2, AVIOTM8, and CHGMTTM. When found to be non-zero, the step adjustment is made to the end timestamp and the amount of time the step was adjusted is added to the job ADJUSTTM data element.

The calculations involved in step adjustment are described by data element below.

PCTCON2/
PCTCON8: This decimal value set in the USER modeling exit is the multiplier of device specific connect time. PCTCON2 is the connect time adjustment factor operating on PGMC3420. PCTCON8 is the connect time factor for PGMC3480. After the increase/decrease amount of time is computed, the end timestamp of the step is decremented by the original device connect time and the altered connect time is applied. Specification of values greater than 1 cause the end timestamp to increase, while values less than 1 cause a decrease in the end timestamp. The connect time for tape drives must be available in the CA MICS variable PGMCTAPE when using this method.

AVIOTM2/
AVIOTM8: This factor should be supplied in decimal seconds and is used to compute a connect time adjustment value that is added to the step end timestamp. Connect time is only available under MVS/XA. Thus, percent connect time adjustment is not possible under MVS/370. Under MVS/370, code in the USER modeling exit must compute a connect time value before any adjustment can be made. An example of this follows:

```

/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step:  IOA Tape Drive Simulation Analysis      |
| Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy |
| Cmd  SAS Statements for:  Specifying USER exit        |
| - -----                                             |
| - %MACRO PERUSEL;                                       |
| -   PGMC3480 = PGME3480 * 0.001;                       |
| -   PGMC3420 = (PGMETAPE - PGME3480 ) * 0.01;         |
| -   PGMCTAPE = PGMC3480 + PGMC3420;                   |
| -   AVIOTM2 = 0.002;                                   |
| - %MEND PERUSEL;                                       |
| ***** BOTTOM OF DATA *****                      |
|                                                         |
|                                                         |
|                                                         |
|                                                         |
\-----/

```

Figure 2-155. User Modeling Exit Panel Example Device Connect Time

When AVIOTM2 or AVIOTM8 is a non-zero quantity, the device connect time is computed based on the number of device type EXCPs. The original device type connect time is subtracted from the end timestamp and the adjusted time is added. AVIOTM2 applies to 3420 devices, while AVIOTM8 applies to 3480 devices.

AVIOTMx and PCTCONx cannot be used for the same step/device type combinations. A warning message is displayed when this condition occurs and no step adjustments are made.

The last adjustment factor is CHGTTM, representing Change Mount Time. The mount time a step sees is not reported by the SMF data. An example showing a current mount time estimate of four minutes being reduced to three minutes is presented in figure 2-156.

```
/----- SAS Statements -----/
| Command ==>                               Scroll ==> CSR
|
| Inquiry Step: IOA Tape Drive Simulation Analysis
|
| Line Cmds: I Insert D Delete R Repeat M Move C Copy
|
| Cmd SAS Statements for: Specifying USER exit
| - -----
| _ - %MACRO PERUSEL;
| _ - CHGMTTM = -60;
| _ - %MEND PERUSEL;
| ***** BOTTOM OF DATA *****
```

Figure 2-156. User Modeling Exit Panel Example - Reducing Mount Time

The Tape Drive Simulation Analysis inquiry will add the nonspecific and specific tape mounts (PGMNSTMT and PGMSTMT), multiply the result by the specified number of seconds supplied via CHGMTTM, and add the result to the end timestamp.

The above data elements: PCTCON2, PCTCON8, AVIOTM2, AVIOTM8, and CHGMTTM can be selectively applied based on other criteria (i.e., job, class, performance group, program, etc.). At the conclusion of all possible step adjustments, the ADJUSTTM data element is recomputed.

2.4.7.4.5 Add Workload Observations Panel

You can use the optional Add Workload Observation panel to model the effect of adding workloads which allocate tape drives.

The following fields are used on the Add Workload Observations panel:

- *Jobname:** The job name you have chosen for this workload. This name must begin with an asterisk (*) to avoid any possible confusion with an actual job being analyzed.
- Start Time:** The time at which this simulated workload begins using tape drives.
- End Time:** The time at which this simulated workload finishes using tape drives.
- Tapes Needed:** The number of drives that this simulated workload will require.
- Warning:** A flag to indicate that a previously defined added workload is no longer within the report hour range (e.g., the report hour range was changed on the main panel).


```

/----- Add Workload Observations -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step: I0A Tape Drive Simulation Analysis |
| Report date:  01DEC87   Report hours:  18 - 23 |
| Line Cmds:  I Insert  D Delete  C Copy  M Move  R Reset |
|
|          Start Time          End Time          Tapes Needed
| Cmd  *Jobname  (hh : mm : ss)  (hh : mm : ss)  (1-99)  Warning
| -  - - - - -  --  --  --  --  --  --  --  --
| -  *          -- : -- : --  -- : -- : --  --  --
| -  *          -- : -- : --  -- : -- : --  --  --
| -  *          -- : -- : --  -- : -- : --  --  --
| -  *          -- : -- : --  -- : -- : --  --  --
| -  *          -- : -- : --  -- : -- : --  --  --
| ***** BOTTOM OF DATA *****
|
\-----/

```

Figure 2-157. Add Workload Observations Panel

2.4.7.4.6 Move Workload Observations Panel

You can use the optional Move Workload Observations panel to model the effect of moving workloads which allocate tape drives to a different time period.

Jobname: The name of the job whose movement to a different time period is to be modeled. This is the name of an actual job whose data is included in the data being analyzed.

Old Start Time: The time (expressed to an accuracy of hundredths of a second) at which this actual job was read into the system. The high level of accuracy for this timestamp is required to ensure that the modeling procedure will be able to match this control statement with the correct job.

New Start Time: The time to which the modeling procedure is to move this job.

```

/----- Move Workload Observations -----\
| Command ==>                               Scroll ==> CSR |
| |
| Inquiry Step:  IOA Tape Drive Simulation Analysis |
| Report date:   01DEC87   Report hours:  18 - 23   |
| |
| Line Cmds:  I Insert  D Delete  C Copy  M Move  R Reset |
| |
|              Old Start Time          New Start Time
| Cmd   Jobname  (hh : mm : ss . th)   (hh : mm : ss . th)
| - - - - - - - - - - - - - - - - - - - - - - - - - - -
| - - - - - - - - - - - - - - - - - - - - - - - - - - -
| - - - - - - - - - - - - - - - - - - - - - - - - - - -
| - - - - - - - - - - - - - - - - - - - - - - - - - - -
| - - - - - - - - - - - - - - - - - - - - - - - - - - -
| - - - - - - - - - - - - - - - - - - - - - - - - - - -
| ***** BOTTOM OF DATA *****
|
|
|
\-----/

```

Figure 2-158. Move Workload Observations Panel

2.4.7.4.7 Alternate Input File Panel

The Alternate Input File panel allows you to specify a SAS file that is not a part of the CA MICS online database as input to the inquiry. You must specify the MVS data set name and the individual SAS data set names.

The optional input file is used to read data from a previously created work file. Creating a work file allows data being selected from the CA MICS database to be saved for later use. Reading this previously saved work file is then possible in a subsequent run of the inquiry.

The Alternate Input File panel requires that all named CA MICS equivalent files be specified. If a named file does not exist, you can code any name so long as there is no file in your SAS library by that name (e.g., NONE). In this case, BAT_ST and BAT_TS may not be needed or are already contained in the alternate BATPGM File (saved selected output from a previous running of this inquiry).

The alternate equivalent CA MICS BATJOB and BAT_JS are specified only if you want to merge these files to obtain access to the JOBCLASS, JOBMCLS, JOBINQTM, or JOBFETTM data elements. These fields can be left blank if they are not needed.

```
/----- Alternate Input File -----\  
| Command ==>  
|  
| Inquiry Step:  IOA Tape Drive Simulation Analysis  
|  
| Specified Time-span:  DETAIL  
|  
| Equivalent CA MICS file(s):  BATPGM    BAT_ST    BAT_TP  
| Alternate SAS file(s)      ==>  _____  
|  
| Merge JOB information:  
|   Equivalent CA MICS file(s):  BATJOB    BAT_JS  
|   Merge SAS files              ==>  _____  
|  
| SAS data set allocation instructions:  
|   Data set name               ==>  _____  
|   Disposition                 ==>  SHR  (SHR/OLD)  
|  
\-----/
```

Figure 2-159. Alternate Input File Panel

2.4.7.4.9 Define Report Description Element Names Panel

You can use the Define report description element names exit to supply additional data elements to be displayed on the Summary or Detail Report. The text should contain a list of data elements whose names can be any data element found in the BATPGM File or data element created in the USER modeling exit. The listed data elements are displayed after the standard and optional data elements.

```
/----- SAS Statements -----  
| Command ==>                               Scroll ==> CSR  
|  
| Inquiry Step: IOA Tape Drive Simulation Analysis  
|  
| Line Cmds: I Insert D Delete R Repeat M Move C Copy  
|  
| Cmd SAS Statements for: Defining report element names  
| - -----  
| _ - %MACRO USERVAR;  
| - -----  
| - -----  
| - -----  
| - -----  
| - -----  
| - -----  
| - -----  
| - -----  
| - -----  
| - -----  
| - -----  
| _ - %MEND USERVAR;  
| ***** BOTTOM OF DATA *****  
|-----/
```

Figure 2-161. Define Report Description Element Names Panel

```
/----- SAS Statements -----/
| Command ==>                               Scroll ==> CSR
|
| Inquiry Step: IOA Tape Drive Simulation Analysis
|
| Line Cmds: I Insert   D Delete   R Repeat   M Move    C Copy
|
| Cmd  SAS Statements for: Defining report element names
| -  -----
| _ - %MACRO USERVAR;
| _ - PERFGRP PGMCPUTM;
| _ - LABEL PERFGRP = 'PERF#GROUP#NUMBER';
| _ - LABEL PGMCPUTM= 'PROGRAM#CPU#TIME'
| _ - %MEND USERVAR;
| ***** BOTTOM OF DATA *****
```

Figure 2-162. Define Report Description Element Names Example

The report formats data element labels on three lines. If user-supplied data elements are added to the report, a user-defined label may also be supplied by ending the data element list with a semicolon and supplying a label statement. The label statement can designate the line on which to start each word in the supplied label with a pound sign (#).

The example shown in Figure 2-162 shows the addition of the elements PERFGRP holding the step's control performance group number and PGMCPUTM holding the step CPU time to a report and associating the labels PERF GROUP NUMBER and PROGRAM CPU TIME with the data elements.

This example would generate a title that appears on the report as:

PERF	PROGRAM
GROUP	CPU
NUMBER	TIME

2.4.8 FICON Channel Performance Analysis Inquiry

The introduction of FICON (fiber optic connectivity) channels into the mainframe platform has helped resolve I/O performance problems by:

- improving the path speed and throughput
- simplifying I/O configuration design by reducing the number of channels needed
- permitting installations to locate devices at remote sites, up to 100km away.

To manage these new channels, IBM has introduced several performance metrics into the SMF type 73 record, which are included in the HARPCA file.

To gain a complete picture of FICON performance, you must be able to view performance in the same manner as the channel itself and to relate those measurements to services times. This requires developing estimates of two new FICON specific metrics: Open Exchanges and Credits.

Open Exchanges (or OEs for short) may be thought of as an MPL for the channel; that is, a measure of concurrency. Credits are a hardware entity, similar to a buffer that contain both the channel program and data for the in-flight I/O operation.

There are specific hardware limits placed on both Open Exchanges and Credits, and exhausting either of these resources will result in channel saturation and poor performance.

This query is designed to compute these new measurements and summarize the data according to the user's specifications and to then produce a series of graphs and reports, depending on execution options.

The following sections describe the facility's operation and reports:

- 1 - Standard Output
- 2 - Technique Tutorial
- 3 - Case Study
- 4 - Component Operation

2.4.8.1 Standard Output

FICON Channel Performance Analysis provides a simple graphic and a series of summary and detail level reports that display both performance and resource usage metrics. Examples of each are included in this section.

The Processor Complex & Channel Configuration report provides a listing of all CECs, LPARs and related FICON channel definitions for the input data. This report can be used to determine the identity of the CECs for later use in data filtering and to better understand which channels are both defined and active on their respective LPARs.

The Channel Service Time graphs are produced for each selected channel and provide a visual analysis of channel service time versus both Open Exchanges and Credits. This permits the analyst to quickly understand the macro-level performance of each channel and to determine if any specific problem areas exist.

The Channel Resource Summary Usage report is used to summarize all performance and resource measurements for each selected channel across the specified time intervals, and additionally to demonstrate the contribution of each LPAR to the overall channel workload. This format allows the analyst to understand the overall performance of each channel and how each LPAR affects the channel performance. A typical use of this report would be to first identify one or more peak times and then execute this report, selecting those time intervals as a follow on to the Channel Service Time graphs discussed above.

The Channel Resource Usage Detail report produces a composite of performance and resource data for each RMF interval for each selected channel. This report can be used to identify periods where a channel is performing poorly.

The sample output below is printed at the beginning of each of the reports and graphs, and depicts the data selection criteria that was specified during query composition. The categories are common to most of the Performance Management product and are not discussed in detail here.

The CEC (Processor Complex) entry refers to the serial number of up to six processor complexes to be selected. The Channels entry is a list of up to eight channels to be selected.

Leaving an entry blank results in a default selection of "ALL", indicating that no filtering is being performed.

Data Criteria Selection Summary:

CA MICS DB Units Specified: P
 Date Range(s) Specified: 01JANyy-31DECyy
 SYSID(s) Selected: XAL1 XAD1
 Zone(s) Selected: 1
 Hour Range(s) Selected: 10-11 14-15
 CEC(s) Selected for Analysis: 0851
 Channels Selected for Analysis: 42 44

The following sections discuss:

- 1 - Processor Complex &
Channel Configuration Report
- 2 - Channel Service Time Graphs
- 3 - Channel Resource Usage Summary Report
- 4 - Channel Resource Usage Detail Report

2.4.8.1.1 Processor Complex & Channel Configuration Report

The Processor Complex & Channel Configuration report provides a listing of all CECs, LPARs and related FICON channel definitions for the input data. This report can be used to determine the identity of the CECs for later use in data filtering and to better understand which channels are both defined and active on their respective LPARs.

CA MICS Performance Manager														1
FICON Channel Performance Analysis														
Processor Complex and Channel Configuration Report														
For Monday, August 30, yyyy														
CEC Serial: 0851 Processor Type: 2084 Model: 307														
LCSS: 0	CA11	CA31	CA71	n/a	n/a	n/a	XAD1	XAE1	XAT1	XA55	XE21	XE44	XE61	n/a n/a
Channels:														
41 FC_S	Y	Y	Y	*	*	*	Y	Y	Y	Y	Y	Y	Y	* *
42 FC_S	Y	Y	Y	*	*	*	Y	Y	Y	Y	Y	Y	Y	* *
43 FC_S	Y	Y	Y	*	*	*	Y	Y	Y	Y	Y	Y	Y	* *
44 FC_S	Y	Y	Y	*	*	*	Y	Y	Y	Y	Y	Y	Y	* *
45 FC_S	N	N	N	*	*	*	N	N	N	N	N	N	N	* *

A 'Y' value indicates that a channel is defined and active.
 A 'N' value indicates that a channel is defined but inactive.
 A '*' value indicates that a channel is undefined

2-86B Processor Complex & Channel Configuration report

The Processor Complex and Channel Configuration report is created by the FICON Channel Performance Analysis query. This report is a listing of all CECs (PR/SM LPAR Central Complex) and their respective LPARs that were located within the input data. Additionally, for each LCSS within each CEC, all FICON channels detected are mapped against their associated LPAR and displayed accordingly.

The fields in the above report include:

CEC Serial	The low-order four bytes of the PRSMCEC data element. This identifies a particular processor complex whose channel data is being analyzed.
LCSS	The Logical Channel Subsystem ID. This one byte field, with a range of values from 0-3, represents a logical grouping of 256 channels. This field is also referred to as the CSSID (Channel Sub-system ID).
Note	On the same report row as LCSS are listed the LPARs being analyzed within the query. Any MVS image executing under VM, or any non-MVS image detected will be excluded, as valid RMF measurement data will not be available.
Processor Type	The model family of the Processor Complex under study. This value is obtained from CA MICS data element CPUMOD (Processor Model Family).
Model	The model description of the Processor Complex, including the model number, derived from data element CPCMODID (CPC Model Identifier).
Channels	A composite of two separate data elements: CHPID (Channel Path ID) and PCAACRNM (Channel Path Acronym). This field will identify each channel being analyzed and also explain the usage of the channel; for example, undefined; switched, or

basic modes.

2.4.8.1.2 Channel Service Time Graphs

The Channel Service Time graphs are produced for each selected channel and provide a visual analysis of channel service time versus both Open Exchanges and Credits. This allows you to quickly understand the macro-level performance of each channel and determine if any specific problem areas exist.




Figure 2-86C shows an example of the Plot of Channel Service Times created by the FICON Channel Performance Analysis query. Two separate plots are created for each channel being analyzed: one for Open Exchanges and one for Credits. In large configurations, such as those with more than 256 channels and using more than one Logical Channel Subsystem, the same channel ID will be used more than once. Therefore, the LCSS identifier (CCSID) will be used to differentiate the channels and will correctly generate multiple plots.

The following measurements are shown in Figure 2-86C:

Service Times (ms)	The Average Service Time is the amount of time used for I/O operations, on average. This is computed as the sum of DVACNN (Device Connect Time), DVADIS (Device Disconnect Time) and DVAPEN (Device pending Time) for all active devices attached to the channel, divided by the total I/O count (SSCH).
Estimated Open exchanges	The Estimated Open Exchanges represent an effective MPL (multi-programming level) for the channel and are the number of I/O operations active concurrently, on average. A complete description of this field and its computation is presented in Section 2.4.8.1.3
Estimated Credits	The Estimated Credits are another measure of channel usage. Credits are an internal hardware buffer used to house a single frame during I/O operations. A complete description of this field and its computation is presented in Section 2.4.8.1.3

2.4.8.1.3 Channel Resource Usage Summary Report

The Channel Resource Summary Usage report is used to summarize all performance and resource measurements for each selected channel across the specified time intervals, and to demonstrate the contribution of each LPAR to the overall channel workload. This format allows you to understand the overall performance of each channel and how each LPAR affects the channel performance. A typical use is to identify one or more peak times using the Channel Service Time graphs and execute this report, selecting those time intervals as a follow up process.

CA MICS Performance Manager														1
FICON Channel Performance Analysis														
Channel Resource Usage Summary Report, by SYSID														
For Monday, August 30, yyyy														
Ending Date: 20JULyyyy Duration: 3:45:02														
CEC Serial: 0851														
LCSS: 0														
CHPID: 42 Description: FICON Switched DCM in Use: Y CHPID Group: 2 Shared: Y														
Channel Usage: DASD CTC														
BUS Usage Count: 116.2M BUS Util. Pct : 13.33 Total Work Units : 419.5M CHPID Pct Busy : 23.87 CHPID I/Os p/sec : 560														
Bytes Read p/Sec: 3.6M Read Path % Busy : 3.01 Bytes Writn p/Sec: 3.2M Write Path % Busy: 2.74														
Avg Burst Size: 12.3K Avg Svc Time (ms): 1.64 Est. Open Exchngs: 0.84 Estimated Credits: 6														
LPAR Usage Stats:														
	Work	CHPID	Total	Bytes	Read	Total	Bytes	Write	Total		Avg	Average	Avg	
SYSID	Units	% Busy	Bytes	Read	Path	Bytes	Written	Path	I/O	I/Os	Burst	Svc Time	Open	Estimd.
			Read	p/sec	% Busy	Written	p/sec	% Busy	Count	p/sec	Size	(ms)	Xchngs	Credits
CA11	15M	0.85	2G	158.8K	0.13	2G	200.1K	0.17	234K	17	20.6K	3.05	0.05	1
CA31	88M	5.05	20G	1.5M	1.26	19G	1.4M	1.23	3M	259	11.4K	0.58	0.12	1
CA71	13M	0.77	1G	94.9K	0.08	1G	134.7K	0.11	322K	23	9.6K	2.52	0.06	0
XAD1	248M	14.11	14G	1.0M	0.90	10G	795.0K	0.66	2M	162	11.5K	3.25	0.50	3
XAE1	6M	0.38	1G	89.1K	0.07	1G	125.7K	0.10	129K	9	22.4K	2.47	0.02	0
XAT1	1M	0.06	91M	6.7K	0.01	1M	85	0.00	22K	1	4.1K	2.62	0.00	0
XA55	922K	0.05	149M	11.0K	0.01	360K	26	0.00	36K	2	4.0K	1.88	0.00	0
XE21	5M	0.32	1G	110.1K	0.09	1G	104.8K	0.09	308K	22	9.4K	0.69	0.01	0
XE44	9M	0.56	2G	197.8K	0.16	1G	108.3K	0.09	181K	13	22.8K	4.45	0.06	1
XE61	14M	0.81	4G	311.7K	0.26	4G	327.9K	0.27	626K	46	13.7K	0.41	0.01	0
2-86D Channel Resource Usage Summary report														

Figure 2-86D shows an example of the Channel Resource Usage Summary report created by the FICON Channel Performance Analysis query. This report is a summary of channel activity for the selected time periods and can be used to review channel performance at a high-level and to set performance objectives for future measurement efforts. A separate report will be produced for each channel selected, and multiple days of input data will generate a separate group of reports for each day. In large configurations, such as those with more than 256 channels and using more than one Logical Channel Subsystem, the same channel ID will be used more than once. Therefore, the LCSS identifier (CCSID) will be used to differentiate the channels and will correctly generate multiple reports.

The following fields are displayed in Figure 2-86D.

Ending Date	The date of the measurement period, derived from the ENDTS element. A separate group of reports is created for each day of input data.
Duration	The total time covered within this reporting period.
CEC Serial	The low-order four bytes of the PRSMCEC data element. This identifies a particular processor complex whose channel data is being analyzed.
LCSS	The Logical Channel Subsystem ID. This is a one byte field, with a range of values from 0-3, and represents a logical grouping of 256 channels. This field is also referred to as the CSSID (Channel Subsystem ID).
CHPID	The identifier of the channel being reported, as a two byte hexadecimal value, with a range of 00-FF.
Description	The description of the channel from the CHANDSCR element. This helps you to better understand the usage of the channel being reported.
DCM in Use	This field indicates whether Dynamic Channel Management is enabled (Y) or not (N).
CHPID Group	The group characteristics of the channel being reported. This value is derived from data

	element PCAGEN (Channel Path Type Generation) and can be used to better understand the operating parameters of the channel. For example, in our sample report the CHPID Group value of 2 indicates that this channel is a FICON Express channel operating at 2GB per second.
Shared	This field indicates whether the channel is being shared among LPARs (Y) or not (N).
Channel Usage	This field explains the usage of active devices assigned to this channel. Possible values are: DASD, TAPE, CTC, COMM and OTHER. Since FICON channels may be assigned to devices with different characteristics and this can impact channel measurement values, it is important to know the exact usage of the channel.
BUS Usage Count	The number of usage cycles used by the PCI BUS to poll for work and transfer data. This information is obtained from the data element PCATOTBC (Total Bus Cycles Count).
BUS Util. Pct.	The utilization of the PCI BUS, expressed as a percentage. Because the BUS is always busy polling for new work, this value can range from 5 to 12 percent even when no actual work is being performed and no data is being transferred. This field is found in data element PCAPCBUS (Pct Channel Busy - Total) and recomputed by the application as required.
Total Work Units	The number of processor cycles used for the transfer of frames. Each frame is 2112 bytes in length, with a 64 byte header and 2048 bytes containing either the channel program or data. This field is obtained from data element PCATOTUC (Channel Work Unit Count - Total).
CHPID Pct. Busy	The percent of (processor) time that the channel is actively transferring commands and data, and is based on the Total Work Units field described above. This field is found in data element PCAPCBSY (Pct Channel Busy - Total) and recomputed by the application as required.
CHPID I/Os p/sec	The number of I/O operations executed by the channel, per second. This is derived as the sum

of I/O operations (SSCH) for all active devices attached to the channel, divided by the measurement interval time.

Note: All data rates presented use an internal SAS format that will display the value in units of K (1,000), M (1,000,000) or G (1,000,000,000) bytes.

Bytes Read p/Sec The total number of bytes read via the channel, per second, formatted and displayed as noted above. This field is a function of data elements PCATOTRU (Read Data Unit Count - Total), PCADUS (Data Unit Size) and the measurement interval.

Read Path % Busy The portion of channel busy used for read operations, expressed as a percentage of the maximum amount of data that can be read concurrently by the channel.

Bytes Written p/Sec The total number of bytes written by the channel, per second, formatted and displayed as noted above. This field is a function of data elements PCATOTWU (Write Data Unit Count - Total), PCADUS (Data Unit Size) and the measurement interval.

Write Path % Busy The portion of channel busy used for write operations, expressed as a percentage of the maximum amount of data that can be written concurrently, by the channel.

Note: FICON channels have separate paths used for reading and writing operations, and perform both operations concurrently and continually. Because of this feature, this query computes usage values for each of the paths independently. Each path must be considered a separate entity for performance and capacity planning purposes.

Avg Burst Size The Average Burst Size is the amount of data (in bytes) transferred during I/O operations, on average. It is computed as the sum of Total Bytes Read and Total Bytes Written, divided by the total number of I/O operations (SSCH) for all active devices attached to the channel.

Avg Svc Time (ms)	The Average Service Time is the amount of time used for I/O operations, on average. This is computed as the sum of DVACNN (Device Connect Time), DVADIS (Device Disconnect Time) and DVAPEN (Device pending Time) for all active devices attached to the channel, divided by the total I/O count (SSCH).
Est. Open Exchngs	The Estimated Open Exchanges represent an effective MPL (multi-programming level) for the channel and are the number of I/O operations active concurrently, on average. This estimated value is computed as the sum of DVACNN (Device Connect Time), DVADIS (Device Disconnect Time) and DVACMRM (Command Response Time) for all active devices attached to the channel, divided by the measurement interval. The device Command Response Time (CMR) time represents, for an I/O request, the time when the first command in the channel program is sent to the device, until the device indicates it has accepted the command, and is a subset of DVAPEN (Device Pending Time). FICON channel architecture currently permits a maximum of 32 Open Exchanges to be used concurrently. A typical value for most applications is from 2 through 5, although higher values have been observed during benchmarks.
Estimated Credits	The Estimated Credits are another measure of channel usage. Credits are an internal hardware buffer used to house a single frame during I/O operations. Currently, FICON architecture allows for a maximum of either 60 or 120 credits to be used, depending on the model of channel. This number is a maximum value for the channel itself. FICON directors contain credits as well. If fewer credits are available within the FICON director, then the maximum number of credits available will be reduced during the protocol negotiation that follows.
LPAR Usage Stats	
SYSID	The CA MICS SYSID value for the reporting LPAR.
Work Units	The number of processor cycles used for the transfer of frames for this LPAR, displayed

using the format discussed above.

CHPID % Busy	The percent of (processor) time that the channel is actively transferring commands and data. This is based on the Work Units field described above. This field is found in data element PCAPCPBY (Pct Channel Busy - LPAR) and recomputed by the query as required.
Total Bytes Read	The total number of bytes read by the channel, from this LPAR, formatted and displayed as noted above. This field is derived from data element PCATOPRU (Read Data Unit Count - LPAR).
Bytes Read p/Sec	The total number of bytes read by the channel, per second, from this LPAR, formatted and displayed as noted above.
Read Path % Busy	The portion of channel busy used for read operations, from this LPAR, expressed as a percentage of the maximum amount of data that can be read concurrently.
Total Bytes Written	The total number of bytes written by the channel, from this LPAR, formatted and displayed as noted above. This field is derived from data element PCATOPWU (Write Data Byte Count - LPAR).
Bytes Writn p/Sec	The total number of bytes written by the channel, per second, from this LPAR, formatted and displayed as noted above.
Write Path % Busy	The portion of channel busy used for write operations, from this LPAR, expressed as a percentage of the maximum amount of data that can be written concurrently.
Total I/O Count	The sum of I/O operations (SSCH) of all active devices attached to this channel for this LPAR.
I/Os p/sec	The number of I/O operations, per second performed by the channel on behalf of devices used by this LPAR. The value is computed the same as the CHPID I/Os p/sec field above.
Avg Burst Size	The Average Burst Size is the amount of data transferred during I/O operations, on average, for devices used by this LPAR. The value is computed the same as the Avg Burst Size field

above.

Avg Svc Time (ms)	The Average Service Time is the amount of time used for I/O operations, on average, for devices used by this LPAR. This value is computed the same as the Avg Svc Time (ms) field above.
Avg Open Xchngs	The number of Estimated Open Exchanges in use for I/O operations performed on behalf of this LPAR. This value is computed in same the manner as the Est. Open Exchngs field above.
Estimd. Credits	The estimated number of credits in use for I/O operations performed on behalf of this LPAR. This value is computed in the same manner as the Estimated Credits field above.

2.4.8.1.4 Channel Resource Usage Detail Report

The Channel Resource Usage Detail report produces a composite of performance and resource data for each RMF interval for each selected FICON channel. This report would be produced and typically used to identify problem periods for a channel deemed to be performing poorly.

CA MICS Performance Manager FICON Channel Performance Analysis Channel Resource Usage Detail Report For Monday, August 30, yyyy															
Ending Date: 20JULyyyy															
CEC Serial: 0851															
LCSS: 0															
CHPID: 42 Description: FICON Switched DCM in Use: Y CHPID Group: 2 Shared: Y															
Device Usage: DASD CTC															
Interval	BUS	BUS	Total		CHPID	CHPID	Bytes	Read	Bytes	Write	Average	Average	Est.		
End	Usage	Util	Work	CHPID	I/O	I/Os	Read	Path	Written	Path	Burst	Service	Open	Est.	
Time	Count	Pct	Units	% Busy	Count	p/sec	p/sec	% Busy	p/sec	% Busy	Size	Time (ms)	Exchns	Credits	
10:00:00	6.7M	11.57	19.2M	16.40	413,577	459	3.3M	2.83	3.7M	3.10	15K	2.05	0.87	7	
10:15:00	8.1M	14.08	30.8M	26.48	699,250	780	2.8M	2.38	3.3M	2.80	7K	1.87	1.34	5	
10:30:00	9.6M	16.61	37.5M	31.93	426,576	473	2.9M	2.47	3.3M	2.82	13K	2.50	1.11	8	
10:45:00	9.8M	16.97	40.5M	34.60	535,592	595	3.7M	3.12	2.9M	2.42	11K	2.28	1.26	8	
11:00:00	10.4M	17.85	43.6M	37.13	503,964	559	4.5M	3.77	2.1M	1.76	11K	2.19	1.14	7	
11:15:00	9.8M	16.86	39.5M	33.74	431,397	479	3.7M	3.12	2.3M	1.93	12K	2.61	1.18	8	
11:30:00	10.1M	17.56	41.3M	35.35	471,276	525	3.9M	3.29	3.3M	2.77	13K	2.92	1.45	10	
11:45:00	9.5M	16.34	38.8M	33.15	523,432	582	2.8M	2.34	3.2M	2.74	10K	2.40	1.31	8	
14:15:00	6.7M	11.56	22.1M	18.89	690,749	767	4.5M	3.76	4.8M	4.05	12K	1.12	0.74	4	
14:30:00	6.1M	10.52	18.5M	15.81	514,899	571	4.4M	3.70	3.8M	3.21	14K	0.79	0.37	3	
14:45:00	6.4M	11.15	20.6M	17.73	531,494	593	4.6M	3.90	4.4M	3.73	15K	0.89	0.44	4	
15:00:00	5.8M	10.12	19.3M	16.44	548,833	608	3.7M	3.12	3.4M	2.85	11K	0.66	0.32	2	
15:15:00	5.3M	9.19	15.0M	12.88	409,228	455	2.4M	2.01	2.4M	2.02	10K	0.80	0.30	2	
15:30:00	5.1M	8.86	12.7M	10.86	332,739	369	2.1M	1.81	2.0M	1.73	11K	0.79	0.24	1	
15:45:00	6.2M	10.71	19.4M	16.60	532,462	592	4.3M	3.60	3.7M	3.10	13K	0.96	0.48	3	
2-86E Channel Resource Usage Detail Report															

Figure 2-86E shows an example of the Channel Resource Usage Detail report created by the FICON Channel Performance Analysis query. This report is a listing of channel activity for the selected time periods by measurement interval, and can be used to review channel performance at a detail level and to set performance objectives for future measurement efforts. A separate report is produced for each channel selected, and multiple days of input data generate a separate group of reports for each day. In large configurations, such as those with more than 256 channels and using more than one Logical Channel Subsystem, the same channel ID will be used more than once. Therefore, the LCSS identifier (CCSID) is used to differentiate the channels and correctly generate multiple reports.

The following fields are displayed in Figure 2-86E:

Ending Date	The date of the measurement period, derived from the ENDTS element. A separate group of reports is created for each day of input data.
Duration	The total time covered within this reporting period.
CEC Serial	The low-order four bytes of the PRSMCEC data element. This identifies a particular processor complex whose channel data is being analyzed.
LCSS	The Logical Channel Subsystem ID. This is a one byte field, with a range of values from 0-3, and represents a logical grouping of 256 channels. This field is also referred to as the CSSID (Channel Subsystem ID).
CHPID	The identifier of the channel being reported, as a two-byte hexadecimal value, with a range of 00-FF.
Description	The description of the channel from the CHANDSCR element. This helps you understand the usage of the channel being reported.
DCM in Use	This field indicates whether Dynamic Channel Management is enabled (Y) or not (N).
CHPID Group	The group characteristics of the channel being reported. This value is derived from data element PCAGEN (Channel Path Type Generation) and can be used to better understand the operating parameters of the channel. For example, in our sample report the CHPID Group value of 2 indicates that this channel is a FICON Express channel operating at 2GB per second.

Shared	This field indicates whether the channel is being shared among LPARs (Y) or not (N).
Channel Usage	This field explains the usage of active devices assigned to this channel. Possible values are: DASD, TAPE, CTC, COMM and OTHER. Since FICON channels may be assigned to devices with different characteristics and this can impact channel measurement values, it is important to know the exact usage of the channel.
Interval End Time	The time of the measurement period, derived from the ENDTS element. A separate line on the report is be created for each measurement interval of input data.
BUS Usage Count	The number of usage cycles used by the hardware PCI BUS to poll for work and transfer data. This information is obtained from the data element PCATOTBC (Total Bus Cycles Count).
BUS Util. Pct	The utilization of the hardware PCI BUS, expressed as a percentage. Because the BUS is always busy polling for new work, this value can range from 5 to 12 percent even when no actual work is being performed and no data is being transferred. This field is found in data element PCAPCBUS (Pct Channel Busy - Total) and recomputed by the application as required.
Total Work Units	The number of processor cycles used for the transfer of frames. Each frame is 2112 bytes in length, with a 64 byte header and 2048 bytes containing either the channel program or data. This field is obtained from data element PCATOTUC (Channel Work Unit Count - Total).
CHPID % Busy	The percent of (processor) time that the channel is actively transferring commands and data, and based on the Total Work Units field described in Figure 2-86E. This field is found in data element PCAPCBSY (Pct Channel Busy - Total) and recomputed by the application as required.

CHPID I/O count	The total number of I/O operations executed by the channel during the measurement interval. This is derived as the sum of I/O operations (SSCH) for all active devices attached to the channel.
CHPID I/Os p/sec	The number of I/O operations executed by the channel, per second. This is derived as the sum of I/O operations (SSCH) for all active devices attached to the channel, divided by the measurement interval time.
Note:	All data rates presented use an internal SAS format that will display the value in units of K (1,000), M (1,000,000) or G (1,000,000,000) bytes.
Bytes Read p/Sec	The total number of bytes read by the channel, per second. This field is a function of data elements PCATOTRU (Read Data Unit Count - Total), PCADUS (Data Unit Size) and the measurement interval.
Read Path % Busy	The portion of channel busy used for read operations, expressed as a percentage of the maximum amount of data that can be read concurrently by the channel.
Bytes Written p/Sec	The total number of bytes written by the channel, per second. This field is a function of data elements PCATOTWU (Write Data Unit Count - Total), PCADUS (Data Unit Size) and the measurement interval.
Write Path % Busy	The portion of channel busy used for write operations, expressed as a percentage of the maximum amount of data that can be written concurrently, by the channel.
Note:	FICON channels have separate paths used for reading and writing operations, and perform both operations concurrently and continually. Because of this feature, this query computes usage values for each of the paths independently. Each path must be considered a separate entity for performance and capacity planning purposes.

Average Burst Size	The Average Burst Size is the amount of data (in bytes) transferred during I/O operations, on average. It is computed as the sum of Total Bytes Read and Total Bytes Written, divided by the total number of I/O operations (SSCH) for all active devices attached to the channel.
Average Service Time (ms)	The Average Service Time is the amount of time used for I/O operations, on average. This is computed as the sum of DVACNN (Device Connect Time), DVADIS (Device Disconnect Time) and DVAPEN (Device pending Time) for all active devices attached to the channel, divided by the total I/O count (SSCH).
Est. Open Exchngs	The Estimated Open Exchanges represent an effective MPL (multi-programming level) for the channel and are the number of I/O operations active concurrently, on average. This estimated value is computed as the sum of DVACNN (Device Connect Time), DVADIS (Device Disconnect Time) and DVACMRTM (Command Response Time) for all active devices attached to the channel, divided by the measurement interval. The device Command Response Time (CMR) represents, for an I/O request, the time when the first command in the channel program is sent to the device, until the device indicates it has accepted the command, and is a subset of DVAPEN (Device Pending Time). FICON channel architecture currently permits a maximum of 32 Open Exchanges to be used concurrently. A typical value for most applications is from 2 through 5, although higher values have been observed during benchmarks.

Est.
Credits

The Estimated Credits are another measure of channel usage. Credits are an internal hardware buffer used to house a single frame during I/O operations. Currently, FICON architecture allows for a maximum of either 60 or 120 credits to be used, depending on the model of channel. This number is a maximum value for the channel itself; FICON directors contain credits as well. If fewer credits are available within the FICON director, then the maximum number the of credits available will be reduced during protocol negotiation that follows.

2.4.8.2 Technique Tutorial

To properly analyze the FICON channel measurements, data from several CA MICS files is required. The main source of measurement data is the CA MICS HARPCA file, which is built from the SMF type 73 record and a basic performance profile can be built from this data alone. However, to fully understand the meaning of the metrics presented in the HARPCA data, one must examine the LCU and device level data as well, and combine it with the channel level measurements. Two new estimated metrics are needed: Open Exchanges, which is a measure of I/O operations concurrency, and Credits, which are a hardware controlled internal buffer that contain both the channel program and data for a given operation. The methodology for obtaining the measurements is presented below.

First, the HARCPU file is read up and the CEC Serial Number is saved, along with some general information about the processor complex. This is done because FICON channels typically operate in a shared mode across LPARs and the basic HARPCA measurement includes a complex level set of measurements along with those of the recording LPAR. To build a complete picture of channel performance, you must be able to group all data for a given processor complex together and the HARCPU file will provide the means to do this.

The main set of metrics is obtained from the HARPCA file. Additionally, beginning with z/OS 1.5 executing on the z990 platform, a new facility was added called the Logical Channel Control Subsystem or LCSS. The purpose of this facility is to permit more than 256 channels to be present on a given platform and to permit LPARs to be mapped onto those channels. Each LCSS can contain up to 256 channels, so that currently, a fully configured z990 system can contain 31 LPARs, using 32 central processors, and mapped against 512 channels. The LCSS is identified by a one byte field called the CSSID, or Channel Subsystem ID and this field resides in the SMF type 73 record. The FICON Performance Analysis application will locate and save the CSSID value to ensure that each FICON channel can be properly mapped against its owning LPAR, and that its measurement data is correctly summarized both at the LPAR and processor complex (CEC) levels.

The Open Exchange metric is an estimate of channel I/O operations concurrency and may be thought of as a channel MPL. It is a function of how busy the channel is during a point in time and is driven by the combination of device

connect and disconnect times. Normally, a channel is not considered "busy" during disconnect times. FICON channels, however, generally do not disconnect during typical operations but rather hold the Open Exchange until the I/O operation has completed. This is done to avoid the overhead of freeing and rebuilding the Open Exchange, given that most disconnect times are due to cache misses, are normally short in duration and hopefully rare. Certain types of I/O operations actually can result in the channel disconnecting and in those instances, the Open Exchange is terminated and freed. However, this is not typically seen in DASD operations.

To compute the (estimated) Open Exchanges, you must first summarize device level measurement data, specifically, the DVACNN (Device Connect Time), DVADIS (Device Disconnect Time) and a new metric, DVACRTM (Command Response Time). The Command Response Time is actually a subset of DVAPEN (Pending Time) and represents the time from the first command being sent to a given device until the device has indicated it has successfully received the command. The Open Exchange computation is essentially a variation of Little's law and is performed by dividing the sum of the above measurements by the duration of the measurements. You can either compute the Open Exchanges based on each device, or total all devices within a Logical Control Unit (LCU) and perform the calculation at that level, or even wait until all measurements have been completed for the channel. Given that this is an estimate, any of the methods may be used with the same degree of accuracy.

Channel-to-Channel devices do a variety of different operations and can seriously impact measurement efforts on FICON channels. CTC devices issue a command called a "Prepare" and then wait for a response. During this time, the device is considered "disconnected" and disconnect time continues to accumulate. This can result in very large values for disconnect time and will subsequently distort the calculation for Open Exchanges. Additionally, CTC devices often use what is called a "Seldom Ending Channel Program," which causes a device to appear "busy." To address this issue, the FICON Performance Analysis application currently will zero out the disconnect time values for any CTC device encountered and compute the Open Exchanges based on device connect and command respond times only.

An additional issue is in regard to non-MVS systems operating on the same mainframe platform. Remember that the SMF type 73 data contains both global (complex) level measurements and

all operating systems residing on the same complex will contribute to these measurements even though only MVS-type systems currently record any LPAR level measurements. Moreover, non-MVS systems will not create any RMF type data, so the sum of all LPAR data may not equal the totals represented by the complex-level measurements. Even MVS systems that operate as "guests" under VM may not record their measurement data completely and can impact FICON channel measurements. To address some of these issues, the FICON Performance Analysis application will delete any measurement data from any MVS system that appears to be executing under VM. It is thought to be better to discard any data that is suspect rather than include it.

Once the device level data has been processed and summarized by LCU, the SCPIOC data is then read up and merged with it. The summarized LCU measurement data is then prorated among all active channels by the IOCPT element, the Number of Times Path Was Taken value. This ensures that any channel not active will be excluded from the prorating process and that the prorated results will properly reflect the channel activity.

The prorated results are merged with the HARPCA data so that a complete picture of channel performance is available. The other estimated metric, Credits, can now be estimated since it requires that an average I/O burst size be first calculated and the result applied to the Open Exchanges. The average I/O burst size is simply the total bytes read and written for the channel, divided by the total Start Subchannel Count. Since the Open Exchanges represent the number of concurrent channel operations, the number of Credits will be a function of this value and the average size (in bytes) of each operation. Again, these are only estimates of these values, and while the methodology presented is believed to be accurate, having these values actually measured and reported through the RMF data would obviously be more effective.

The complete measurement results are presented in a variety of graphs and reports.

2.4.8.3 Case Study

To be added later.

2.4.8.4 Component Operation

This Component Operation Section describes the online inquiry panels required for the operation of FICON Channel Performance Analysis routines.

Entries in the online inquiry panels define the types of reports to be generated. In particular, these panels define:

- o The cycles from the CA MICS HARCPU, HARPCA, SCPIOC and HARDVA files to be processed.
- o The type of report(s) to be produced.

The FICON Channel Performance Analysis reports provided are:

- o Processor Complex & Channel Configuration Report - Describes the configuration under study, in terms of Processor Complexes (CECs), LPARs and Channels.
- o Channel Resource Usage Summary Report - Reports the resources used and performance observed for each channel selected for study, subsetting the measurement data by LPAR.
- o Channel Resource Usage Detail Report - Reports the individual resource usage and performance observed for each selected channel by RMF measurement interval.

Additionally, graphs of both Open Exchanges and Credits, together with Channel Service Times can be optionally produced, subject to the data selection criteria below.

The following options are common to both control panels:

- Cycle(s): One or more cycles of the CA MICS files that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a cycle from the CA MICS database will usually contain a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided by the DATE and HOUR options.
- Database: The single-character CA MICS identifier for the database to be accessed. Valid values are the database unit identifiers available at your

installation. Contact your CA MICS System Administrator if you do not know what database units are available, or enter a question mark (?) to view a list of your database units. You can select up to five databases per inquiry.

Allow Run-Time Override:	The indication of whether the database identifier and/or the Cycles specification can be overridden at inquiry execution time.
Date:	Up to two date ranges (DDMONYY format) to be used for analysis. This option is normally used to drop the extra day from a cycle that is due to the timing of the SMF dump.
Hour:	One to three ranges of hours for which data is to be reported. For an individual hour, leave the end hour blank or make it the same as the start hour. You can specify a value from 00 to 23.
Zone:	The specification that zone selection is to be applied. If zone 1 definition is 8:00 a.m. to 6:00 p.m., you can specify reports covering that time period by entering a 1 in the ZONE field.
Sysid:	The specification that SYSIDs are to be selected for reporting. You can enter up to six SYSIDs to be selected. If you leave the field blank, no testing of SYSIDs will be done.
CEC Serial:	The specification that processor complexes are to be selected for reporting. You can enter up to six CEC Serial Numbers to be selected. If you leave the field blank, no testing of processor complexes will be done.
Channels:	The specification that FICON channels are to be selected for reporting. You can enter up to eight channel id numbers to be selected. The channel ids are entered as 2 digit hexadecimal values, with a range of 00 to ff. If you leave the field blank, no testing of channels will be done.

The operation of the zone and time selection options depends on definitions made during the installation of the CA MICS system. If you are not sure about the ZONE definitions for your installation or the CA MICS database update procedures,

consult your CA MICS System Administrator.

FICON Channel Performance Analysis	
Command ==>	+
Modifying CA MICS Inquiry: SAMPLE - FICON Sample Cycle(s) ==> 01 - 01	
Data base ==> P (FORALL)	
Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N)	
Data Selection: Allow run-time override	
Date ==> 01JANyy - 31DECyy	==> N (Y/N)
Hour ==> 10 - 11 14 - 15	==> N (Y/N)
Zone ==> 1	==> N (Y/N)
SYSID ==> SYS1 SYS1	==> N (Y/N)
CEC Serial ==> 0801 9205	
Specify the channels you wish to report on (Leave blank to select all)	
Channels ==> 42 44	
Reporting Options:	
Generate CEC & Channel Configuration: YES (YES/NO)	
Generate Service Time Summary Plot: YES (YES/NO)	
Generate Resource Summary Usage: YES (YES/NO)	
Generate Resource Detail Usage: YES (YES/NO)	

2.5 Virtual Storage Management Analysis

This section introduces virtual storage concepts to form a background for discussing the management of virtual storage in MVS/XA and MVS/ESA. The composition of MVS virtual storage is described, with the goal of defining the methods that MVS uses to manage virtual storage and the measurements available for analyzing it.

Virtual storage control is explained for MVS/XA only, because RMF does not provide any measurement of virtual storage for MVS/370.

The Virtual Storage Management Analysis inquiries enable you to track the use of virtual storage by MVS/XA at any supported time-span and summarization pattern. Reports and plots display system measurements of virtual storage allocation. More detailed reports of Common Area storage usage are also available, including allocation information within storage subpools and protection keys. In addition to system level and Common Area reporting, you can display virtual storage resource distribution patterns using either plots or reports for any Private Area being measured by RMF.

Section 2.5.4, Virtual Storage Concepts in MVS/XA, defines the terms, explains the basic dynamics of many of the elements used in virtual storage analysis, and describes the types of virtual storage measurements provided by RMF. Depending on your familiarity with these terms and concepts, you may find it useful to read this section before proceeding with Sections 2.5.2 and 2.5.3.

The following sections describe the operation of the Virtual Storage Management Analysis inquiries and how to use the reports that are produced:

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output
- 4 - Virtual Storage Concepts in MVS/XA
- 5 - Case Study
- 6 - Component Operation
- 7 - Requirements

2.5.1 Functional Description

The Virtual Storage Management Analysis inquiries allow you to track the systemwide utilization of virtual storage, Common Area virtual storage usage, and virtual storage usage by workload in MVS/XA. Any supported CA MICS timespan and summarization pattern can be implemented.

The System Storage Usage Report and Common Area Storage Usage Report use data from the MVS Virtual Storage System Usage File (SCPVSM). The System Storage Usage Plot obtains data from the same source. The Private Storage Usage Report and Plot use data from the MVS Virtual Storage Private Usage File (SCPVSU).

You specify the timespan desired during inquiry definition. The range of data that can be reported from a CA MICS database depends on the timespans that were defined for these files and the number of cycles defined in each timespan.

For the Private Storage Usage Report and Plot to produce meaningful results, the application to be analyzed must have been reported upon by RMF.

2.5.2 Usage Guidelines

This section discusses the application of the concepts discussed in Section 2.5.4, Virtual Storage Concepts in MVS/XA. In particular, the following topics are discussed:

- 1 - Horizontal Versus Vertical Address Expansion
- 2 - Major Impacts Upon Virtual Storage Usage
- 3 - Monitoring Virtual Storage Changes

2.5.2.1 Horizontal Versus Vertical Address Expansion

HORIZONTAL EXPANSION OF VIRTUAL ADDRESSING

A major historical theme in the design of MVS has been what can be viewed as a horizontal expansion of the virtual storage space.

In order to provide virtual storage constraint relief (VSCR), applications which compete for virtual storage might be run on physically separate processors. For instance, IMS, which tends to require a large Common System Area (CSA) and CICS, which needs a large private region (PVT), might be run on separate processors. This separation of a large CSA and a large Private Area onto separate processors would constitute expansion of these two workloads into two separate virtual address ranges in the two processors.

The increased capability to spread the virtual addressing range horizontally within a single MVS system has been implemented in MVS/370 through the use of cross-memory address spaces. Such horizontal expansion has permitted major functions to be isolated in their own address spaces, thus effectively expanding the virtual addressing range available beyond the 16MB range of MVS/370 and the 2 GB range of MVS/XA by allowing access by one address space into another one, thus logically including the address space identification (ASID) as a high-level portion of the address.

For the MVS/XA operating system itself, this horizontal expansion has been implemented via supplemental address spaces:

- o PCAUTH
- o TRACE
- o GRS
- o CATALOG
- o CONSOLE
- o ALLOCAS
- o LLA
- o SMF

JES3, with its use of an auxiliary address space is an example of a subsystem making use of cross-memory services to provide this same sort of horizontal expansion of virtual storage.

This trend is being continued with the advent of PR/SM and MVS/ESA.

VERTICAL EXPANSION OF VIRTUAL ADDRESSING

The advent of MVS/370 provided a vertical expansion of the virtual storage addressing range, since it removed the former restriction that all private regions share one 16MB address range. The subsequent expansion of this range to 2 GB with the arrival of MVS/XA continued this trend. However, programs cannot take advantage of the virtual addresses above 16 MB without at least some modification and reassembly. Thus, since many of the larger users of virtual storage are subsystems like JES3, VTAM, or IMS, it is often necessary to upgrade to a new release of such a product to obtain needed VSCR.

2.5.2.2 Major Impacts Upon Virtual Storage Usage

SYSTEM QUEUE AREA (SQA)

SQA size is specified at IPL via the member IEASYSnn in SYS1.PARMLIB and does not change thereafter. If it is overspecified, virtual address space will be wasted. If it is underspecified, requests for unavailable SQA will be satisfied from the CSA with little or no performance impact.

COMMON SYSTEM AREA (CSA)

Like SQA, CSA is specified at IPL via the member IEASYSnn and does not change thereafter. The size specified for CSA is subtracted from the bottom PLPA address. If this does not end on a megabyte boundary, it is rounded down to the next lower megabyte boundary to determine the top of the Private Area. Thus, small changes in the allocated size of CSA can cause the Private Area to decrease by one megabyte if one of these boundaries is crossed.

Major users of CSA are IMS, VTAM, and JES3. A detailed knowledge of these subsystems and the parameters controlling their generation and operation is necessary to assist in managing their use of virtual storage.

PLPA, FLPA, AND MLPA

Analysis of virtual storage usage by the PLPA, FLPA, and MLPA depends heavily upon detailed analysis of the modules found in these areas. You should be certain that non-IBM modules which do not have to be in these areas for performance reasons are excluded. There should be periodic checking for obsolete versions of system modules, such as the JES2 load module HASPSSSM. Searching for module names like HASPSSSX and TSTXSSSM can aid in this sort of housecleaning.

PRIVATE AREA

The management of virtual storage in the Private Area depends upon a detailed knowledge of the subsystem or application under concern. For instance, for JES2, specifying inappropriate sizes for buffers can lead to wasted, unusable space in the JES2 Private Area.

2.5.2.3 Monitoring Virtual Storage Changes

Monitoring changes in system and private virtual storage areas over an extended period of months or weeks can assist you in determining the need for system software upgrades which take advantage of various VSCR techniques. Such monitoring can also warn you of changes in the sizes of areas like PLPA.

You can analyze virtual storage allocations in the various system areas with the System Storage Usage Reports and Plots described in Sections 2.5.3.2 and 2.5.3.5. Additional detailed information is available for the Common System Area (CSA) with reports described in Section 2.5.3.3. The methods for producing these reports and plots are described in Sections 2.5.6.1, 2.5.6.2, and 2.5.6.4.

You can monitor the virtual storage allocations in any particular address space with the Private Storage Usage Reports and Plots described in Sections 2.5.3.4 and 2.5.3.6. Methods for producing these reports and plots are described in Sections 2.5.6.3. and 2.5.6.5

2.5.3 Standard Output

The VSM storage analysis reports give you the capability to track the allocation of virtual storage at any supported time-span and summarization level. Results of the analyses are presented in tabular, plot, and graphic output formats.

Virtual storage requirements at both the system and private levels are reported to aid you in evaluating the impact and the growth of existing applications. You can use the selection parameters described in Section 2.5.6, Component Operation, to limit reporting to specific systems, times, users, and output options.

Described below are the report formats of the System Storage Usage Reports and Plots, Common Area Usage Reports, and Private Storage Usage Reports and Plots. Common title and report elements are also defined.

- 1 - Common Report Elements
- 2 - System Storage Usage Report
- 3 - Common Area Storage Usage Report
- 4 - Private Storage Usage Report
- 5 - System Storage Usage Plot
- 6 - Private Storage Usage Plot

2.5.3.1 Common Report Elements

This section describes elements common to all VSM analysis reports. Options that are unique to a specific report are documented with that report.

TITLE ELEMENTS

Each report has five lines of title information available. All reports give the name of the report, the name of the component, and the run date. Additional headings detail the data selected, the summary level of the report, and the time periods reported by column. For the monthly report, the column title is month/year; for the weekly report, the first day of the week is given; for the daily report, the date is given as month/day/year; for the detail report, the column title is time interval-day. For the hourly level, the hour is given as the column title and the day of the month is given in the report title.

The following elements are in the standard title portion of each report:

Installation Name:	On some reports, the installation name that is defined at CA MICS installation time. This value is centered in the first line of the report. For more information about this value, consult your CA MICS System Administrator.
Component Name:	The name of the performance component, centered in the first or second line of each page.
Report Name:	The name of the report, centered in the third line of the heading. It cannot be modified.
SYSID or System:	The CA MICS logical SYSID being reported. Data from multiple systems cannot be combined using standard control parameters.
ZONE:	If data is being reported by zone, the zone number is shown.
RUNDATE:	The date the report was run.

ADDITIONAL ELEMENTS

Reporting Interval:	The first row of some reports shows the date or date/time combination of the period being

reported. If the date being reported on is from the WEEKS timespan, the date reflects the starting date of the week, as derived from the CA MICS variable STARTTS. In all other cases, the date reflects the ending date, derived from the CA MICS variable ENDTS. In monthly reports, only the year and month portions of the date are displayed.

% CHANGE: The last column on some reports shows the percent change for each report item (row) between the base period (first data column) and the last period (last data column). If data for only one period is reported, then dashes (---) are displayed. If the base period is zero, then N/A is printed to indicate that a percent change could not be calculated.

2.5.3.2 System Storage Usage Report

The System Storage Usage Report tracks the allocation of virtual storage on a monthly, weekly, daily, hourly, or detail level. Figure 2-165 shows an example of a monthly System Storage Usage Report and Figure 2-166 shows the hourly form of the report. In the hourly form, data for one day is displayed by hour using the DAYS time-span of the MVS Virtual Storage System Usage File (SCPVSM).

REPORT OPTIONS

The System Storage Usage Report supports standard time selection options, as well as SYSID selection and external data selection. These options are controlled by the control options discussed in Section 2.5.6.

DATA ELEMENTS

The elements reported are divided into the following categories:

- o Average Map
- o CSA Average Values
- o SQA Average Values
- o Key Average Values

All categories contain above and below 16MB usage values.

AVERAGE MAP

Average Map Usage is shown in terms of storage allocated to the MVS nucleus, the System Queue Area, the Common Storage Area, the Pageable Link Pack Area, and the Private Area.

NUCLEUS:	The size of the nucleus area, including the nucleus load module and nucleus control blocks. These values come from VSMANUCL and VSMANUCH elements.
CSA:	The storage assigned to the Common Storage Area. This value is contained in the VSMACSAL and VSMACSAH elements.
LPA:	The storage assigned to the Pageable Link Pack Area from the VSMALPAL and VSMALPAH elements. Also reported separately are the FLPA (Fixed LPA) storage from the VSMAFLPL and VSMAFLPH

elements, and the MLPA (Modified LPA) storage from the VSMAMLPL and VSMMLPH elements.

SQA: The storage assigned to the System Queue Area from the VSMASQAL and VSMASQAH elements. Because SQA is totally fixed storage, the size of the SQA also indicates the number of real storage frames used.

PRIVATE: The storage that is assigned to the Private Areas for all MVS address spaces. These are from the VSMAPVTL and VSMAPVTH elements.

CSA AVERAGE VALUES

CSA Average Values are shown in terms of the total allocated storage and the average amount of unused storage.

FREE PAGES (BYTES): The total amount of storage available in the CSA, from the VSMAFCSL and VSMAFCSH elements.

LARGEST FREE BLOCK: The total amount of storage available in the CSA in a contiguous block, from the VSMABCSL and VSMABCSH elements.

SQA AVERAGE VALUES

SQA Average Values are shown in terms of the total allocated storage and the average amount of unused storage.

FREE PAGES (BYTES): The total amount of storage available in the SQA, from the VSMAFSQL and VSMAFSQH elements.

LARGEST FREE BLOCK: The total amount of storage available in the SQA in a contiguous block, from the VSMABSQL and VSMABSQH elements.

KEY AVERAGE VALUES

Key Average Values are useful in tracking exceptional conditions occurring with the system storage.

SQA EXPANSION INTO CSA: The total amount of storage used when the SQA required expansion. This storage is obtained from the CSA; such an occurrence generally represents a performance degradation. These values are found in the VSMACSQL and VSMACSQH elements.

PLPA PACKING LOSS:	The total amount of storage that is allocated but not used among load modules in the PLPA. Such a loss occurs because all PLPA storage is allocated in pages, but most often load modules do not fill a page completely. These values are found in the VSMAPAKL and VSMAPAKH elements.
PLPA REDUNDANT SPACE:	The total amount of storage that is allocated redundantly because modules exist in PLPA and also in either the FLPA or MLPA. These values are found in the VSMADUPL and VSMADUPH elements.
MAXIMUM USER REGION:	The largest amount of storage that can be specified as a REGION parameter on a JCL statement. These values are found in the VSMAUSRL and VSMAUSRH elements.

Additionally, it is possible optionally to produce all of the above-reported elements with MINIMUM/MAXIMUM values to provide for more detailed analysis.

MONTHLY SYSTEM STORAGE USAGE REPORT							
SYSTEM = CPU2		RUNDATE = JUL 5, 1988					
MONTHS (YY/MM):	87-12	88-01	88-02	88-03	88-04	88-05	CHANGE %
AVERAGE MAP - ABOVE 16 MEG							
PRIVATE	2009M	2009M	2004M	2001M	2001M	2001M	-0.42 %
CSA	7998K	8548K	12M	16M	16M	16M	106.12 %
MLPA	44K	44K	44K	44K	44K	44K	---
FLPA	16K	16K	16K	16K	16K	16K	---
PLPA	3968K	3968K	4096K	4220K	4220K	4220K	6.35 %
SQA	8512K	8512K	8512K	8512K	8512K	8512K	---
NUCLEUS	2464K	2464K	2464K	2464K	2464K	2464K	---
AVERAGE MAP - BELOW 16 MEG							
NUCLEUS	412K	412K	409K	418K	424K	424K	2.91 %
SQA	640K	640K	640K	640K	640K	640K	---
PLPA	2541K	2520K	2406K	2306K	2312K	2312K	-9.01 %
FLPA	4K	4K	4K	4K	4K	4K	---
MLPA	4K	5K	4K	4K	4K	4K	---
CSA	6638K	6658K	6775K	6866K	6856K	6856K	3.26 %
PRIVATE	6140K	6140K	6140K	6140K	6140K	6140K	---
PREFIX SAVE AREA	4K	4K	4K	4K	4K	4K	---
CSA AVG VALUES - ABOVE 16M							
FREE PAGES (BYTES)	3776K	3643K	9864K	13M	13M	13M	268.19 %
LARGEST FREE BLOCK	3684K	3441K	9524K	12M	12M	12M	251.03 %
ALLOCATED AREA SIZE	4319K	5219K	3057K	3996K	3507K	3554K	-17.70 %
CSA AVG VALUES - BELOW 16M							
FREE PAGES (BYTES)	1222K	1131K	1549K	1816K	1739K	1674K	36.94 %
LARGEST FREE BLOCK	1065K	966K	1406K	1691K	1572K	1536K	44.24 %
ALLOCATED AREA SIZE	5646K	5715K	5712K	5171K	5280K	5809K	2.89 %

Figure 2-165. Sample Monthly System Storage Usage Report (Part 1 of 2)

MONTHLY SYSTEM STORAGE USAGE REPORT							
SYSTEM = CPU2		RUNDATE = JUL 5, 1988					
MONTHS (YY/MM):	87-12	88-01	88-02	88-03	88-04	88-05	CHANGE %

SQA AVG VALUES - ABOVE 16M							

FREE PAGES (BYTES)	6476K	6495K	6578K	6518K	6530K	6590K	1.74 %
LARGEST FREE BLOCK	6405K	6447K	6516K	6457K	6462K	6545K	2.19 %
ALLOCATED AREA SIZE	8512K	8512K	8512K	8512K	8512K	8512K	---

SQA AVG VALUES - BELOW 16M							

FREE PAGES (BYTES)	13K	8K	11K	6K	9K	9K	-31.21 %
LARGEST FREE BLOCK	5K	4K	4K	4K	4K	4K	-26.98 %
ALLOCATED AREA SIZE	638K	638K	638K	638K	638K	638K	-0.10 %

KEY AVG VALUES - ABOVE 16M							

SQA EXPANSION TO CSA	0K	0K	0K	0K	0K	0K	N/A
PLPA PACKING LOSS	9K	9K	10K	11K	11K	10K	15.49 %
PLPA REDUNDANT SPACE	13K	13K	13K	13K	13K	13K	---
MAXIMUM USER REGION	2000M	2000M	1996M	1992M	1992M	1992M	-0.42 %

KEY AVG VALUES - BELOW 16M							

SQA EXPANSION TO CSA	67K	83K	76K	100K	95K	88K	31.05 %
PLPA PACKING LOSS	6K	5K	6K	6K	6K	6K	-2.5.8 %
PLPA REDUNDANT SPACE	3K	4K	2K	2K	2K	2K	-12.82 %
MAXIMUM USER REGION	5988K	5988K	5988K	5988K	5988K	5988K	---

Figure 2-165. Sample Monthly System Storage Usage Report (Part 2 of 2)

HOURLY SYSTEM STORAGE USAGE REPORT SYSTEM = MVS1 RUNDTE = JUL 5, 1988										
HOUR/DATE (HR-MM/DD):	08-06/22	09-06/22	10-06/22	11-06/22	12-06/22	13-06/22	14-06/22	15-06/22	16-06/22	CHANGE %
AVERAGE MAP - ABOVE 16 MEG										
PRIVATE	2001M	2001M	2001M	2001M	2001M	2001M	2001M	2001M	2001M	---
CSA	16M	16M	16M	16M	16M	16M	16M	16M	16M	---
MLPA	44K	44K	44K	44K	44K	44K	44K	44K	44K	---
FLPA	16K	16K	16K	16K	16K	16K	16K	16K	16K	---
PLPA	4248K	4248K	4248K	4248K	4248K	4248K	4248K	4248K	4248K	---
SQA	8512K	8512K	8512K	8512K	8512K	8512K	8512K	8512K	8512K	---
NUCLEUS	2464K	2464K	2464K	2464K	2464K	2464K	2464K	2464K	2464K	---
AVERAGE MAP - BELOW 16 MEG										
NUCLEUS	424K	424K	424K	424K	424K	424K	424K	424K	424K	---
SQA	640K	640K	640K	640K	640K	640K	640K	640K	640K	---
PLPA	2312K	2312K	2312K	2312K	2312K	2312K	2312K	2312K	2312K	---
FLPA	4K	4K	4K	4K	4K	4K	4K	4K	4K	---
MLPA	4K	4K	4K	4K	4K	4K	4K	4K	4K	---
CSA	5832K	5832K	5832K	5832K	5832K	5832K	5832K	5832K	5832K	---
PRIVATE	7164K	7164K	7164K	7164K	7164K	7164K	7164K	7164K	7164K	---
PREFIX SAVE AREA	4K	4K	4K	4K	4K	4K	4K	4K	4K	---
CSA AVG VALUES - ABOVE 16M										
FREE PAGES (BYTES)	12M	12M	12M	12M	12M	12M	12M	12M	12M	-1.46 %
LARGEST FREE BLOCK	12M	12M	12M	12M	12M	12M	12M	12M	12M	-1.58 %
ALLOCATED AREA SIZE	3969K	4076K	4141K	4133K	4128K	4128K	4134K	4162K	4168K	5.00 %
CSA AVG VALUES - BELOW 16M										
FREE PAGES (BYTES)	1893K	1884K	2137K	1960K	1873K	1862K	1877K	1868K	1867K	-1.36 %
LARGEST FREE BLOCK	1838K	1840K	1833K	1529K	1544K	1544K	1544K	1526K	1544K	-16.01 %
ALLOCATED AREA SIZE	3993K	3991K	3998K	4302K	4288K	4288K	4288K	4305K	4288K	7.37 %

Figure 2-166. Sample Hourly System Storage Usage Report (Part 1 of 2)

HOURLY SYSTEM STORAGE USAGE REPORT SYSTEM = MVS1 RUNDTE = JUL 5, 1988										
HOUR/DATE (HR-MM/DD):	08-06/22	09-06/22	10-06/22	11-06/22	12-06/22	13-06/22	14-06/22	15-06/22	16-06/22	CHANGE %

SQA AVG VALUES - ABOVE 16M										

FREE PAGES (BYTES)	7480K	7437K	7420K	7416K	7420K	7417K	7421K	7454K	7445K	-0.46 %
LARGEST FREE BLOCK	7480K	7437K	7420K	7416K	7420K	7417K	7419K	7419K	7421K	-0.77 %
ALLOCATED AREA SIZE	8512K	8512K	8512K	8512K	8512K	8512K	8512K	8512K	8512K	---

SQA AVG VALUES - BELOW 16M										

FREE PAGES (BYTES)	256K	253K	246K	245K	241K	238K	235K	237K	235K	-7.89 %
LARGEST FREE BLOCK	256K	253K	243K	241K	236K	235K	233K	235K	235K	-7.89 %
ALLOCATED AREA SIZE	384K	386K	396K	398K	404K	404K	406K	404K	404K	5.26 %

KEY AVG VALUES - ABOVE 16M										

SQA EXPANSION TO CSA	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
PLPA PACKING LOSS	9K	9K	9K	9K	9K	9K	9K	9K	9K	---
PLPA REDUNDANT SPACE	13K	13K	13K	13K	13K	13K	13K	13K	13K	---
MAXIMUM USER REGION	1992M	1992M	1992M	1992M	1992M	1992M	1992M	1992M	1992M	---

KEY AVG VALUES - BELOW 16M										

SQA EXPANSION TO CSA	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
PLPA PACKING LOSS	6K	6K	6K	6K	6K	6K	6K	6K	6K	---
PLPA REDUNDANT SPACE	2K	2K	2K	2K	2K	2K	2K	2K	2K	---
MAXIMUM USER REGION	7012K	7012K	7012K	7012K	7012K	7012K	7012K	7012K	7012K	---

Figure 2-166. Sample Hourly System Storage Usage Report (Part 2 of 2)

Area and of selected Private Areas. The following sections include discussions of each of these topics:

- 1 - Virtual Storage Map for MVS/XA
- 2 - Initialization and Management of Virtual Storage
- 3 - VSM Management and Measurement Macros
- 4 - Virtual Storage Subpools
- 5 - Storage Protection
- 6 - Virtual Storage Measurement by RMF
- 7 - Allocation Measurements and Their Interpretation

2.5.3.3 Common Area Storage Usage Report

The Common Area Storage Usage Report tracks the allocation of virtual storage at a monthly, weekly, daily, hourly, or detail level. Figure 2-167 shows an example of a monthly Common Area Storage Usage Report and Figure 2-168 shows the hourly form of the report. In the hourly form, data for one day is displayed by hour using the DAYS time-span of the MVS Virtual Storage System Usage File (SCPVSM).

REPORT OPTIONS

The Common Area Storage Usage Report supports standard time selection options, as well as SYSID selection and external data selection. These options are controlled by the control options discussed in Section 2.5.6.

DATA ELEMENTS

The elements reported are divided into the following categories:

- o CSA Average Values
- o SQA Average Values
- o CSA Average by Key

All categories contain above and below 16MB usage values.

CSA AVERAGE VALUES

CSA Average Values is shown in terms of the total allocated storage and the average amount of unused storage.

ALLOCATED CSA: The amount of storage (pages) allocated actively in the CSA, from the VSMAUCSL and VSMAUCSH elements.

FREE PAGES (BYTES): The total amount of storage available in the CSA, from the VSMAFCSL and VSMAFCSH elements.

LARGEST FREE BLOCK: The total amount of storage available in the CSA in a contiguous block, from the VSMABCSL and VSMABCSH elements.

ALLOCATED AREA SIZE: The total amount of storage (pages) allocated in the CSA, from the VSMATCSL and VSMATCSH elements. This measurement represents a "high water" mark of storage allocated within the CSA. The difference between this measurement and the ALLOCATED CSA represents storage fragmentation.

SQA AVERAGE VALUES

SQA Average Values are shown in terms of the total allocated storage and the average amount of unused storage.

ALLOCATED SQA: The amount of storage (pages) allocated actively in the SQA, from the VSMAUSQL and VSMAUSQH elements.

FREE PAGES (BYTES): The total amount of storage available in the SQA, from the VSMAFSQL and VSMAFSQH elements.

LARGEST FREE BLOCK: The total amount of storage available in the SQA in a contiguous block, from the VSMABSQL and VSMABSQLH elements.

ALLOCATED AREA SIZE: The total amount of storage (pages) allocated in the SQA, from the VSMATSQL and VSMATSQLH elements. This measurement represents a "high water" mark of storage allocated within the SQA. The difference between this measurement and the ALLOCATED CSA represents storage fragmentation for SQA storage residing below 16MB. For SQA residing above 16MB, it normally has an allocated area equal to the specified ESQA size.

CSA AVERAGE BY KEY

CSA Average by Key values are shown in terms of the total allocated storage and represent MVS Storage Keys 0 through 7 and 8 through F.

KEY x The amount of storage (pages) allocated in MVS
ALLOCATION: Storage Key x in the CSA, from the VSMAxCSL and
 VSMAxCSH elements, where x represents one of the
 MVS storage protection keys.

A reporting option permits either reporting by storage key or reporting by storage subpool and key. Additionally, it is possible optionally to produce all of the above-reported elements with MINIMUM/MAXIMUM values to provide for more detailed analysis.

MONTHLY COMMON AREA STORAGE USAGE REPORT BY KEY							
SYSTEM = CPU2 RUNDATE = JUL 5, 1988							
MONTHS (YY/MM):	87-12	88-01	88-02	88-03	88-04	88-05	CHANGE %
CSA AVG VALUES - > 16M							
ALLOCATED CSA	4222K	4904K	2716K	2416K	2412K	2583K	-32.5.81 %
FREE PAGES (BYTES)	3776K	3643K	9864K	13M	13M	13M	262.5.19 %
LARGEST FREE BLOCK	3684K	3441K	9524K	12M	12M	12M	251.03 %
ALLOCATED AREA SIZE	4319K	5219K	3057K	3996K	3507K	3554K	-17.70 %
CSA AVG VALUES - < 16M							
ALLOCATED CSA	5348K	5443K	5149K	4949K	5020K	5093K	-4.77 %
FREE PAGES (BYTES)	1222K	1131K	1549K	1816K	1739K	1674K	36.94 %
LARGEST FREE BLOCK	1065K	966K	1406K	1691K	1572K	1536K	44.24 %
ALLOCATED AREA SIZE	5646K	5715K	5712K	5171K	5280K	5809K	2.89 %
SQA AVG VALUES - > 16M							
ALLOCATED SQA	2035K	2016K	1933K	1993K	1981K	1921K	-5.56 %
FREE PAGES (BYTES)	6476K	6495K	6578K	6518K	6530K	6590K	1.74 %
LARGEST FREE BLOCK	6405K	6447K	6516K	6457K	6462K	6545K	2.19 %
ALLOCATED AREA SIZE	8512K	8512K	8512K	8512K	8512K	8512K	---
SQA EXPANSION TO CSA	0K	0K	0K	0K	0K	0K	N/A
SQA AVG VALUES - < 16M							
ALLOCATED SQA	693K	714K	705K	733K	726K	719K	3.65 %
FREE PAGES (BYTES)	13K	8K	11K	6K	9K	9K	-31.21 %
LARGEST FREE BLOCK	5K	4K	4K	4K	4K	4K	-26.98 %
ALLOCATED AREA SIZE	638K	638K	638K	638K	638K	638K	-0.10 %
SQA EXPANSION TO CSA	67K	83K	76K	100K	95K	88K	31.05 %

Figure 2-167. Sample Monthly Common Area Storage Usage Report (Part 1 of 2)

MONTHLY COMMON AREA STORAGE USAGE REPORT BY SUBPOOL							
SYSTEM = CPU2 RUNDATE = JUL 5, 1988							
MONTHS (YY/MM):	87-12	88-01	88-02	88-03	88-04	88-05	CHANGE %
CSA AVERAGE BY KEY - > 16M							
KEY 0 ALLOCATION	527K	570K	541K	593K	603K	583K	10.57 %
KEY 1 ALLOCATION	2916K	3431K	953K	32K	32K	32K	-92.5.90 %
KEY 2 ALLOCATION	0K	0K	0K	0K	0K	0K	N/A
KEY 3 ALLOCATION	57K	59K	57K	64K	63K	63K	11.45 %
KEY 4 ALLOCATION	0K	0K	0K	0K	0K	0K	N/A
KEY 5 ALLOCATION	57K	56K	57K	56K	55K	55K	-2.15 %
KEY 6 ALLOCATION	0K	0K	338K	649K	616K	608K	N/A
KEY 7 ALLOCATION	608K	709K	706K	927K	971K	1183K	94.38 %
KEYS 8-F ALLOCATION	55K	77K	62K	93K	69K	56K	1.86 %
CSA AVERAGE BY KEY - < 16M							
KEY 0 ALLOCATION	1619K	1626K	1629K	1644K	1632K	1615K	-0.25 %
KEY 1 ALLOCATION	171K	153K	150K	158K	154K	153K	-10.72 %
KEY 2 ALLOCATION	0K	0K	0K	0K	0K	0K	N/A
KEY 3 ALLOCATION	99K	116K	94K	110K	109K	117K	17.99 %
KEY 4 ALLOCATION	0K	0K	0K	0K	0K	0K	N/A
KEY 5 ALLOCATION	22K	24K	23K	24K	21K	21K	-4.34 %
KEY 6 ALLOCATION	622K	677K	409K	180K	168K	167K	-73.12 %
KEY 7 ALLOCATION	2683K	2715K	2719K	2702K	2813K	2895K	7.90 %
KEYS 8-F ALLOCATION	130K	129K	122K	128K	121K	123K	-5.30 %

Figure 2-167. Sample Monthly Common Area Storage Usage Report (Part 2 of 2)

HOURLY COMMON AREA STORAGE USAGE REPORT BY KEY										
SYSTEM = MVS1 RUNDTE = JUL 5, 1988										
HOURLY DATE (HR-MM/DD):	08-06/22	09-06/22	10-06/22	11-06/22	12-06/22	13-06/22	14-06/22	15-06/22	16-06/22	CHANGE %
CSA AVG VALUES - > 16M										
ALLOCATED CSA	3965K	4064K	3783K	3584K	3920K	4028K	4109K	4142K	4148K	4.62 %
FREE PAGES (BYTES)	12M	12M	12M	12M	12M	12M	12M	12M	12M	-1.46 %
LARGEST FREE BLOCK	12M	12M	12M	12M	12M	12M	12M	12M	12M	-1.58 %
ALLOCATED AREA SIZE	3969K	4076K	4141K	4133K	4128K	4128K	4134K	4162K	4168K	5.00 %
CSA AVG VALUES - < 16M										
ALLOCATED CSA	3938K	3947K	3694K	3872K	3958K	3969K	3954K	3963K	3964K	0.65 %
FREE PAGES (BYTES)	1893K	1884K	2137K	1960K	1873K	1862K	1877K	1868K	1867K	-1.36 %
LARGEST FREE BLOCK	1838K	1840K	1833K	1529K	1544K	1544K	1544K	1526K	1544K	-16.01 %
ALLOCATED AREA SIZE	3993K	3991K	3998K	4302K	4288K	4288K	4288K	4305K	4288K	7.37 %
SQA AVG VALUES - > 16M										
ALLOCATED SQA	1032K	1074K	1091K	1095K	1091K	1094K	1090K	1057K	1066K	3.37 %
FREE PAGES (BYTES)	7480K	7437K	7420K	7416K	7420K	7417K	7421K	7454K	7445K	-0.46 %
LARGEST FREE BLOCK	7480K	7437K	7420K	7416K	7420K	7417K	7419K	7419K	7421K	-0.77 %
ALLOCATED AREA SIZE	8512K	8512K	8512K	8512K	8512K	8512K	8512K	8512K	8512K	---
SQA EXPANSION TO CSA	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
SQA AVG VALUES - < 16M										
ALLOCATED SQA	384K	386K	393K	394K	398K	401K	404K	402K	404K	5.26 %
FREE PAGES (BYTES)	256K	253K	246K	245K	241K	238K	235K	237K	235K	-7.89 %
LARGEST FREE BLOCK	256K	253K	243K	241K	236K	235K	233K	235K	235K	-7.89 %
ALLOCATED AREA SIZE	384K	386K	396K	398K	404K	404K	406K	404K	404K	5.26 %
SQA EXPANSION TO CSA	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A

Figure 2-168. Sample Hourly Common Area Storage Usage Report (Part 1 of 2)

HOURLY COMMON AREA STORAGE USAGE REPORT BY KEY										
SYSTEM = MVS1 RUNDTE = JUL 5, 1988										
HOUR/DATE (HR-MM/DD):	08-06/22	09-06/22	10-06/22	11-06/22	12-06/22	13-06/22	14-06/22	15-06/22	16-06/22	CHANGE %
CSA AVERAGE BY KEY - > 16M										
KEY 0 ALLOCATION	380K	380K	361K	360K	400K	397K	396K	396K	396K	4.21 %
KEY 1 ALLOCATION	32K	32K	32K	32K	32K	32K	32K	32K	32K	---
KEY 2 ALLOCATION	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
KEY 3 ALLOCATION	268K	310K	181K	97K	261K	309K	352K	363K	364K	35.66 %
KEY 4 ALLOCATION	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
KEY 5 ALLOCATION	52K	52K	52K	52K	52K	52K	52K	52K	52K	---
KEY 6 ALLOCATION	464K	521K	511K	326K	438K	501K	541K	552K	552K	12.5.86 %
KEY 7 ALLOCATION	2760K	2760K	2636K	2710K	2728K	2728K	2728K	2738K	2744K	-0.57 %
KEYS 8-F ALLOCATION	8K	8K	9K	5K	8K	8K	8K	8K	8K	---
CSA AVERAGE BY KEY - < 16M										
KEY 0 ALLOCATION	965K	968K	969K	986K	1006K	1013K	1006K	1006K	1010K	4.70 %
KEY 1 ALLOCATION	104K	104K	104K	104K	104K	104K	104K	104K	104K	0.90 %
KEY 2 ALLOCATION	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
KEY 3 ALLOCATION	96K	90K	91K	135K	149K	144K	146K	151K	147K	53.47 %
KEY 4 ALLOCATION	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
KEY 5 ALLOCATION	28K	28K	18K	17K	28K	28K	28K	28K	28K	---
KEY 6 ALLOCATION	172K	171K	201K	309K	116K	116K	116K	116K	117K	-31.65 %
KEY 7 ALLOCATION	2537K	2549K	2276K	2289K	2520K	2526K	2516K	2516K	2515K	-0.84 %
KEYS 8-F ALLOCATION	36K	36K	32K	28K	34K	36K	36K	39K	40K	11.11 %

Figure 2-168. Sample Hourly Common Area Storage Usage Report (Part 2 of 2)

2.5.3.4 Private Storage Usage Report

The Private Storage Usage Report tracks the use of virtual storage at a monthly, weekly, daily, hourly, or detail level.

REPORT OPTIONS

The Private Storage Usage Report supports standard time selection options as well SYSID selection and external data selection. The main selection criteria will likely be the JOB name. Normally, one or more reports will be produced for each job specified. These options are discussed in Section 2.5.6.

This report can be produced in two formats:

- o OVERVIEW
- o TRENDING

OVERVIEW FORMAT

The OVERVIEW format presents a series of variables dealing with allocated and free private storage, both above and below 16MB. Figure 2-169 shows the hourly form of the report, using the OVERVIEW format. In the hourly form of the report, data from one day is displayed by hour using the DAYS time-span of the MVS Virtual Storage Private Usage File (SCPVSU).

OVERVIEW FORMAT DATA ELEMENTS

The elements reported are divided into the following categories:

- o Job Statistics
- o Average Region Size Values
- o Average Map
- o Average Values
- o Average Subpool Alloc < 16M

With the exception of Average Subpool Allocations, all categories include usage values above and below 16MB.

JOB STATISTICS

Job Statistics include both the number of SAMPLES and INTERVLS used in the report.

SAMPLES: The number of RMF samples used in the input data.

INTERVLS: The number of RMF recording intervals used in the data.

AVERAGE REGION SIZE VALUES

Average Region Size Values include measurements related directly to the user region size and pertain to the amount of storage requested, but not necessarily used.

JCL REQUESTED: The amount of storage requested via the REGION parameter in the JCL. This value is derived from the VSUAREGR element.

ASSIGNED
>< 16M: The amount of storage assigned via exits. These values are found in the VSUAURAH and VSUAURAL elements.

AVERAGE MAP

LSQA/SWA/
229/230: The amount of storage allocated for LSQA and SWA subpools. These values are found in the VSUAALSH and VSUAALSL elements.

FREE LSQA/
SWA/229-30: The amount of storage in the largest free block of the LSQA and SWA subpools. These values are found in the VSUABLSH and VSUABLSL elements.

FREE USER
REGION: The amount of storage in the largest free block of the user region subpools. These values are found in the VSUABUSH and VSUABUSL elements.

ALLOC.
USER
REGION: The amount of storage allocated for the user region subpools. These values are found in the VSUAAUSH and VSUAAUSL elements.

SYSTEM
REGION: The amount of storage allocated in each user region for system (overhead) functions.

AVERAGE VALUES

FREE BYTES
LSQA-230: The amount of storage of free pages of LSQA and SWA subpools. These values are contained in the VSUAFLSH and VSUAFLSL elements.

GETMAIN
BYTES: The largest amount of storage that can be obtained via a GETMAIN macro. These values are found in the VSUAGMLH and VSUAGMLL elements.

AVERAGE SUBPOOL ALLOC. < 16M

Average Subpool Alloc. < 16M includes the measurements of the amount of virtual memory allocated to various subpools below the 16MB line.

SP xxx: The amount of storage allocated to subpool x.
 These values are found in the VSUAxxxL elements,
 where xxx is one of the following: 000, 229,
 230, 236, 237, 251, 252, and 255.

Additionally, it is possible to optionally produce all of the above-reported elements with MINIMUM/MAXIMUM values to provide for detailed analysis. This option is only applicable to reports using the OVERVIEW format, however.

HOURLY PRIVATE STORAGE USAGE REPORT BY JOB										
SYSTEM = TS04 RUNDATE = JUL 1, 1988										
JOB REPORTED ON = JESA										
HOURLY DATE (HR-MM/DD):	08-06/28	09-06/28	10-06/28	11-06/28	12-06/28	13-06/28	14-06/28	15-06/28	16-06/28	CHANGE %

JOB STATISTICS										

SAMPLES IN PERIOD	358.00	359.00	359.00	356.00	360.00	360.00	356.00	361.00	359.00	0.27 %
INTERVALS IN PERIOD	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	---

AVERAGE REGION SIZE VALUES										

JCL REQUESTED	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
ASSIGNED > 16M	2015M	2015M	2015M	2015M	2015M	2015M	2015M	2015M	2015M	---
ASSIGNED < 16M	8172K	8172K	8172K	8172K	8172K	8172K	8172K	8172K	8172K	---

AVERAGE MAP - ABOVE 16 MEG										

LSQA/SWA/229/230	8840K	8840K	8840K	8840K	8840K	8840K	8840K	8840K	8840K	0.00 %
FREE LSQA/SWA/229-30	2005M	2005M	2005M	2005M	2005M	2005M	2005M	2005M	2005M	---
FREE USER REGION	604K	604K	604K	604K	604K	604K	604K	604K	604K	---
ALLOC USER REGION	604K	604K	604K	604K	604K	604K	604K	604K	604K	---

AVERAGE MAP - BELOW 16 MEG										

LSQA/SWA/229/230	165K	165K	165K	164K	164K	165K	165K	168K	168K	1.97 %
FREE LSQA/SWA/229-30	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
FREE USER REGION	5708K	5708K	5708K	5708K	5708K	5708K	5708K	5723K	5727K	0.34 %
ALLOC USER REGION	2184K	2184K	2184K	2184K	2184K	2184K	2184K	2184K	2184K	-0.00 %
SYSTEM REGION	16K	16K	16K	16K	16K	16K	16K	16K	16K	---

Figure 2-169. Sample Hourly Private Storage Usage Report - OVERVIEW Option (Part 1 of 2)

HOURLY PRIVATE STORAGE USAGE REPORT BY JOB										
SYSTEM = TS04		RUNDATE = JUL 1, 1988		JOB REPORTED ON = JESA						
HOURLY DATE (HR-MM/DD):	08-06/28	09-06/28	10-06/28	11-06/28	12-06/28	13-06/28	14-06/28	15-06/28	16-06/28	CHANGE %
AVERAGE VALUES - ABOVE 16M										
FREE BYTES LSQA-230	2005M	2005M	2005M	2005M	2005M	2005M	2005M	2005M	2005M	---
GETMAIN BYTES	2048M	2048M	2048M	2048M	2048M	2048M	2048M	2048M	2048M	---
AVERAGE VALUES - BELOW 16M										
FREE BYTES LSQA-230	0K	0K	0K	0K	0K	0K	0K	0K	0K	N/A
GETMAIN BYTES	8192K	8192K	8192K	8192K	8192K	8192K	8192K	8192K	8192K	---
AVERAGE SUBPOOL ALLOC <16M										
SP 0	1636K	1636K	1636K	1636K	1636K	1636K	1636K	1636K	1636K	-0.00 %
SP 229	25K	24K	25K	24K	24K	24K	25K	25K	24K	-2.06 %
SP 230	68K	68K	68K	68K	68K	68K	68K	71K	72K	5.70 %
SP 236 (SWA)	48K	48K	48K	48K	48K	48K	48K	48K	48K	---
SP 237 (SWA)	4K	4K	4K	4K	4K	4K	4K	4K	4K	---
SP 251	420K	420K	420K	420K	420K	420K	420K	420K	420K	---
SP 252 (REENT.)	48K	48K	48K	48K	48K	48K	48K	48K	48K	0.02 %
SP 255 (LSQA)	20K	20K	20K	20K	20K	20K	20K	20K	20K	-0.49 %

Figure 2-169. Sample Hourly Private Storage Usage Report - OVERVIEW Option (Part 2 of 2)

MONTHLY PRIVATE STORAGE USAGE REPORT - TRENDING OPTION							
SYSTEM = TS04		RUNDATE = JUL 1, 1988		REPORTED ELEMENT = VSUAAUSL: Avg K Alc Pg User Region < 16MB			
MONTHS (YY/MM):	87-12	88-01	88-02	88-03	88-04	88-05	CHANGE %
VIRTUAL STORAGE USE :							
JESA	2176K	2176K	2176K	2177K	2183K	2183K	0.32 %

Figure 2-170. Sample Monthly Private Storage Usage Report - TRENDING Option.

TRENDING FORMAT

The TRENDING format permits reporting on any numeric element present in the selected input file. Figure 2-170 shows an example of a monthly Private Storage Usage Report using the TRENDING format.

2.5.3.5 System Storage Usage Plot

The System Storage Usage Plot tracks the allocation of virtual storage at a monthly, weekly, daily, hourly, or detail level. Figure 2-171 shows an example of a daily System Storage Usage Plot. The plots obtain their data exclusively from the SCPVSM File.

PLOT OPTIONS

The System Storage Usage Plot supports standard time selection options, as well as SYSID selection and external data selection. These options are controlled by the control options discussed in Section 2.5.6.

Additionally, the plot may be produced in two formats:

- o TOTAL
- o BELOW 16 MEG. ONLY

In the TOTAL format, the plot displays all of the measurement variables, both above and below the 16MB. line. The TOTAL option is the default. In the BELOW 16 MEG. ONLY format, only the measurement variables from below 16MB. are displayed. This format will produce a more detailed display of this lower portion of virtual storage than that of the TOTAL option.

DATA ELEMENTS

PREFIX:	The size of the Prefix Save Area, which is fixed at 4K. This is an area used by the hardware and addressed explicitly by the operating system.
PRIVATE:	The storage that is assigned to the private areas for all MVS address spaces. These values are from the VSMAPVTL and VSMAPVTH elements.
CSA:	The storage assigned to the Common Storage Area. This value is contained in the VSMACSAH and VSMACSAH elements.
LPA:	The storage assigned to the Pageable Link Pack Area, from the VSMALPAL and VSMALPAH elements. Also plotted separately are the FLPA (Fixed LPA) storage, from the VSMAFLPL and VSMAFLPH elements, and the MLPA (Modified LPA) storage, from the VSMAMLPL and VSMMLPH elements.

SQA: The storage assigned to the System Queue Area from the VSMASQAL and VSMASQAH elements. Because SQA is totally fixed storage, the size of the SQA also indicates the number of real storage frames used.

NUCLEUS: The size of the nucleus area, including the nucleus load module and nucleus control blocks. This value comes from the VSMANUCL and VSMANUCH elements.




Figure 2-171. Sample Daily System Storage Usage Plot

2.5.3.6 Private Storage Usage Plot

The Private Storage Usage Plot tracks the use of virtual storage at a monthly, weekly, daily, hourly, or detail level.

PLOT OPTIONS

The Private Storage Usage Plot supports standard time selection options as well SYSID selection and external data selection. The main selection criteria will likely be the JOB name. Normally, one or more plots will be produced for each job specified. These options are discussed in Section 2.5.6.

This plot can be produced in two formats:

- o OVERVIEW
- o TRENDING

OVERVIEW FORMAT

The OVERVIEW format presents a series of variables dealing with allocated and free private storage both above and below 16MB. Figure 2-172 shows the daily form of the plot, using the OVERVIEW format.

OVERVIEW FORMAT DATA ELEMENTS

The OVERVIEW format data elements are plotted using the first character of its data element description. Thus, SYSTEM REGION is plotted using S.

SYSTEM REGION:	The amount of storage allocated in each user region for system functions.
ALLOC. USER REGION:	The amount of storage allocated for the user region subpools. These values are found in the VSUAAUSH and VSUAAUSL elements.
FREE USER REGION:	The amount of storage in the largest free block of the user region subpools. These values are found in the VSUABUSH and VSUABUSL elements.
LSQA/SWA/229/230:	The amount of storage allocated for LSQA and SWA subpools. These values are found in the VSUAALSH and VSUAALSL elements.
COMMON	The amount of storage allocated for all Common

AREA: Area uses, including PLPA, CSA, and SQA, as well as their counterparts above 16MB. The Common Area values are computed by the software from the differences between the private area size(s) and the size of addressable storage.




Figure 2-172. Sample Daily Private Storage Usage Plot - OVERVIEW Option




Figure 2-173. Sample Weekly Private Storage Usage Plot - TRENDING Option.

TRENDING OPTION

The TRENDING format permits plotting on any numeric element present in the selected input file. Figure 2-173 shows an example of a weekly Private Storage Usage Plot using the TRENDING format.

2.5.4 Virtual Storage Concepts in MVS/XA

The methodology section is an introduction to the concepts in MVS/XA underlying Virtual Storage Manager (VSM) control of available virtual storage and RMF measurement of its use. Virtual storage is allocated in 4KB increments in either the Common or Private Areas in response to requests for GETMAINS for various subpools of storage. These subpools characterize whether the memory will be page-fixed or fetch-protected, in what major area of virtual storage it will exist, and the direction of allocation. RMF can measure the allocation and fragmentation of these portions of the Common Area and of selected Private Areas. The following sections include discussions of each of these topics:


- 1 - Virtual Storage Map for MVS/XA
- 2 - Initialization and Management of Virtual Storage
- 3 - VSM Management and Measurement Macros
- 4 - Virtual Storage Subpools
- 5 - Storage Protection
- 6 - Virtual Storage Measurement by RMF
- 7 - Allocation Measurements and Their Interpretation

2.5.4.1 Virtual Storage Map for MVS/XA

One of the goals of a virtual system is to relieve application programs of memory management concerns. In a virtual memory system, a program is given a large amount of virtual memory to use. The operating system then allocates real storage to jobs as required. Figure 2-174 shows a simplified view of virtual storage arrangement in an MVS/XA system running several programs. To each program, it appears as if the entire virtual storage of 2 billion bytes (2 gigabytes or 2 GB) is available to it (minus areas used by the system control program).

To provide the necessary virtual storage constraint relief, allowing users to write applications that are substantially larger than those possible in MVS/370 systems, MVS/XA expanded the size of the address to 31 bits. Programs can now address up to 2 GB of virtual storage. To maintain compatibility with MVS/370 and yet provide expanded addressability, MVS/XA provides areas of virtual storage that are defined below the 16MB boundary line and "extended" portions of most of these same areas that are defined above the 16MB boundary line.

Figure 2-174 shows conceptually how this is done. MVS/XA treats both the portion below the 16MB line and the extended portion above the 16MB line as one logical area. In the figure, each job, started task, or TSO user running in an MVS/XA system would have its own copy of the Private Area and the Extended Private Area.







Figure 2-174. Virtual Address Map - MVS/XA

ADDRESSING MODE AND MODULE LOCATION

A user may determine whether any particular program will use 24 bit or 31 bit addressing. The AMODE attribute which is specified to the assembler or linkage editor specifies 31 bit addressing. The default if AMODE attribute is not specified is 24 bit addressing.

A user may also determine whether modules are to be loaded into the User Region (below the 16MB line) or the Extended User Region (above 16MB line) via the RMODE attribute for the assembler or linkage editor. The default if the RMODE attribute is not specified is to load the module into the User Region below the 16MB line. This allows programs written for MVS/370 to run in MVS/XA with generally no modification.

A more detailed map of the different areas within the 2GB address range of virtual storage is shown in Figure 2-175.



EMLPA - Extended MLPA		
EFLPA - Extended FLPA		
EPLPA - Extended PLPA		
ESQA - Extended SQA		
ENUC - Extended Nucleus		
+=====+		16MB Boundary Line
Common Area		
NUC - Nucleus		
SQA - System Queue Area		
PLPA - Pageable Link Pack Area		
FLPA - Fixed Link Pack Area		
MLPA - Modified Link Pack Area		
CSA - Common System Area		
+-----+		
Private Area		
LSQA - Local System Queue Area		
Authorized Region Subpools		
SWA - Scheduler Work Area		
AUK - Authorized User Key(229/230)		
- - - - -		
Free Space		
Free Space		
- - - - -		
User Region Subpools		
- - - - -		20K
System Region		
+-----+		4K
Common Area		
PSA - Prefixed Storage Area		
+-----+		0

Figure 2-175. Detailed Virtual Address Map - MVS/XA

This map shows further detail within each of the four major areas of MVS/XA virtual storage: Extended Private Area, Extended Common Area, Common Area, and Private Area. MVS/XA manages this virtual storage via the Virtual Storage Manager (VSM).

Note that the PSA and the System Region do not have counterparts above the 16MB boundary line. Similarly, the VSM and RSM Work Areas and the Private Area Page Tables have no counterparts below the 16MB boundary line.

2.5.4.2 Initialization and Management of Virtual Storage

Some of these areas of virtual storage are established at the time MVS is initialized and are not subject to any subsequent dynamic management by VSM.

At the time of IPL (Initial Program Load), the IPL control module creates the necessary virtual environment for other programs to be loaded from SYS1.NUCLEUS. The IPL control module and the IRIMs (IPL Resource Initialization Modules) that it subsequently loads into memory load the nucleus, including the PSA and the nucleus map, build the page frame table above the nucleus map, allocate minimal SQA and ESQA areas, and allocate LSQA for the master scheduler. Control is then passed to the NIP (Nucleus Initialization Program) to complete the initialization of various system resources, including the system virtual storage areas.

The NIP, together with some of the RIMs (Resource Initialization Modules) which it invokes, expands SQA and ESQA to match the parameters found in SYS1.PARMLIB. It then creates or resets PLPA and EPLPA (depending on whether the person doing the IPL requested a cold start or a warm start). Next, the modules to be placed in FLPA, EFLPA, MLPA, and EMLPA are loaded. This is followed by the allocation of CSA and ECSA according to the parameters in SYS1.PARMLIB. At this time the NIP has essentially finished establishing the systemwide areas of virtual storage. It then creates various system address spaces, the first of which is the Master Scheduler, to which it finally passes control.

All of these areas which are established by these initialization processes have their sizes set during initialization. The establishment of their overall sizes determines the sizes of the private and extended private areas, since they are the portions of virtual memory below and above the 16MB line that are left over after all the system areas have been designated.

Only four of these system areas are subject to continuing dynamic management by VSM: the SQA, ESQA, CSA, and ECSA areas. The dynamics of virtual memory usage within these areas and within certain critical address spaces in the Private Area can become a major concern, since exhausting available virtual memory in these areas can lead to performance degradation, abnormal termination of these critical address spaces, or, under extreme conditions, the abnormal termination of the entire system.

2.5.4.3 VSM Management and Measurement Macros

There are ten macro instructions which can be used either to request that VSM perform some management function or some measurement function upon virtual storage. These are the GETMAIN, FREEMAIN, CPOOL, VSMLIST, VSMLOC, and VSMREGN macros, plus four others whose functions have been replaced in MVS/XA by the more efficient CPOOL macro. Although these latter four macros are still supported for compatibility, we will not discuss them here.

VSMLOC can be used to verify that a control block or storage location has been allocated. This verification function is useful in operations such as recovery procedures, but is of limited value in the management of virtual storage, and thus will not be discussed further here. The major features of the other macros are summarized below.

GETMAIN AND FREEMAIN

The GETMAIN and FREEMAIN macros form a logical pair. GETMAIN's function is to allocate one or more areas of virtual storage to an active address space. When requesting GETMAIN services, the length of the portion of virtual storage desired, together with its required subpool number, need to be specified. Subpools will be described in more detail later. Their use is a technique used by VSM to classify the memory it is managing. For most subpools, the storage obtained via GETMAIN is released either when the associated task issues a FREEMAIN or at task termination. The VSMLOC macro is a related macro that can be used to verify that a particular portion of virtual storage has been previously allocated via a GETMAIN.

CPOOL

The CPOOL macro performs a set of functions that are superficially similar to those performed by GETMAIN and FREEMAIN. The CPOOL macro performs the following "cell pool services": BUILD, GET, FREE, and DELETE. The purpose of these services is to manage pools of equal sized cells of virtual storage. These are commonly used to construct chains of control blocks or other collections of equal-sized packets of storage.

The BUILD function creates a cell pool. The following information is used by the CPOOL macro when specifying the BUILD function: size of each cell, primary and secondary cell count, subpool number, storage protect key, the address

of the TCB with which this pool is to be associated, and some header text for diagnostic purposes. The primary and secondary cell counts are treated much like primary and secondary space values during disk file allocation. The BUILD process will create as many cells as are specified in the primary cell count. If subsequent requests for cells depletes a cell pool, then the secondary cell count will be used to extend the cell pool into secondary extents.

The GET function of the CP00L macro obtains a cell from the previously built pool. If the pool is depleted, the macro call can specify either that the request fail or that the pool be extended by acquiring enough virtual storage to build the number of additional cells specified in the secondary cell count during the initial building function for the cell pool.

The FREE function of the CP00L macro returns a cell to the pool, whereas the DELETE function deletes the entire cell pool and frees the storage of all the extents of the pool and its associated control blocks.

VSMREGN AND VSMLIST

Two other analytical macros are a part of the VSM-related services available to programs. The first of these, VSMREGN, is used to obtain the address and length of the private area region for a given TCB, both below and above the 16MB line.

The second of these macros, VSMLIST, is used extensively by RMF in its tracking of virtual storage usage. For various portions of virtual storage, this macro can provide several different detailed lists of virtual storage areas:

- o Allocated pages - The addresses and lengths of groups of contiguous pages which have had some or all of their memory allocated via GETMAINS. These measurements will be in multiples of 4KB.
- o Unallocated pages - The addresses and lengths of groups of contiguous pages in which no GETMAINS have allocated space. These measurements will also be in multiples of 4KB.
- o Free space - The addresses and lengths of fragments of memory which, although within allocated pages, have not been reserved for use via GETMAINS. These measurements will not necessarily be in multiples of 4KB.

The ways that RMF makes use of this VSMLIST macro will be explained in more detail later during the discussion of RMF measurement of virtual storage usage.

2.5.4.4 Virtual Storage Subpools

Virtual storage is subdivided into different areas, which are shown in Figure 2-175. Within these areas, virtual storage is further subdivided into categories called subpools.

GETMAIN/FREEMAIN processing of the VSM will accept requests for the subpool numbers shown in the Figure 2-176. The VSM maps these requests into its own somewhat smaller group of subpool numbers and allocates sufficient virtual address space, in increments of 4KB pages, to satisfy any particular GETMAIN request. The VSM will allocate virtual memory within an area either from the bottom upward, or from the top downward, depending upon the subpool involved. Storage allocated in these subpools may be page fixed or pageable, and fetch protected or not fetch protected. All these characteristics are summarized for the various subpools in the Figure 2-176.

Requested Subpool	VSM Subpool	Page Fixed	Fetch Pro- tected	Type of Storage And Allocation Direction	Owner	Subpool FREEMAIN	Automatic Freeing
000-127	000-127	No	Yes	PVT/EPVT/Up	Current Task	Yes	Task
226	226	Yes	No	SQA/Down	None	No	No
227	227	Yes	Yes	CSA/ECSA/Down	None	No	No
228	228	Yes	No	CSA/ECSA/Down	None	No	No
229	229	No	Yes	PVT/EPVT/Down	Current Task	Yes	Task
230	230	No	No	PVT/EPVT/Down	Current Task	Yes	Task
231	231	No	Yes	CSA/ECSA/Down	None	No	No
233	255	Yes	No	LSQA/ELSQA/Down	Current Task	Yes	Task
234	255	Yes	No	LSQA/ELSQA/Down	Jobstep Task	No	Jobstep
235	255	Yes	No	LSQA/ELSQA/Down	Memory	No	No
236	236	No	No	PVT/SWA/Down	Current Task	Yes	Task
237	237	No	No	PVT/SWA/Down	Current Task	Yes	Task
239	239	Yes	Yes	SQA/ESQA/Down	None	No	No
240	000	No	Yes	PVT/EPVT/Up	Current Task	Yes	Jobstep
241	241	No	No	CSA/ECSA/Down	None	No	No
245	245	Yes	No	SQA/ESQA/Down	None	No	No
250	000	No	Yes	PVT/EPVT/Up	Current Task	Yes	Jobstep
251	251	No	Yes	PVT/EPVT/Up	Jobstep Task	Yes	Jobstep
252	252	No	No	PVT/EPVT/Up	Jobstep Task	Yes	Jobstep
253	255	Yes	No	LSQA/ELSQA/Down	Current Task	Yes	Task
254	255	Yes	No	LSQA/ELSQA/Down	Jobstep Task	No	jobstep
255	255	Yes	No	LSQA/ELSQA/Down	Memory	No	No

Figure 2-176. Virtual Storage Subpool Characteristics
PVT/EPVT/SWA - PRIVATE AREA, EXTENDED PRIVATE AREA, AND
SCHEDULER WORK AREA

Subpool 240 has been retained for compatibility with MFT and OS/VS1. This subpool and subpool 250 are translated by VSM into GETMAINS for subpool 0. Requests for subpool 0 by tasks running with storage protection key 0 in supervisor state are translated to subpool 252. Any subpool in a user region will be given the storage protection key of the task that issued the first GETMAIN for that subpool. That task is considered the owner of that subpool and all other requests for that subpool will be given the same storage protection key.

Authorized region subpools consist of the Authorized User Key subpools (AUK), 229 and 230, and the Scheduler Work Area subpools (SWA), 236 and 237. The AUK supports multiple storage protection keys, whereas the entire SWA is assigned storage protection keys of 1.

LSQA/ELSQA - LOCAL SYSTEM QUEUE AREA AND EXTENDED LOCAL
SYSTEM QUEUE AREA

VSM defines subpools 233, 234, and 235 to be equivalent to 253, 254, and 255, respectively. Internally, all are managed as subpool 255. All will have a storage protection key of 0.

CSA/ECSA - COMMON SYSTEM AREA AND EXTENDED COMMON SYSTEM AREA

Subpools 227, 228, 231, and 241 give every possible combination of page fixing and fetch protection. In order to help ensure data and program integrity, VSM supports multiple storage protection keys for the CSA.

SQA/ESQA - SYSTEM QUEUE AREA AND EXTENDED SYSTEM QUEUE AREA

All virtual storage assigned to SQA which is not currently allocated to subpools 226 or 239 is considered to be part of subpool 245. All these SQA subpools are assigned a storage protection key of 0.

Note that subpool 226 is never allocated above the 16MB boundary line.

2.5.4.5 Storage Protection

There are three mechanisms used for protecting memory from unauthorized fetching and storing: page protection, low address protection, and key-controlled protection. Only one of these, key controlled protection, relates to the areas of memory that are actively managed by VSM. However, we will briefly summarize the other two mechanisms.

PAGE PROTECTION

The page protection mechanism controls against improper storing by any reference using a virtual address. It is managed by the use of the page-protection bit which exists for each page in its corresponding page table entry. This type of protection is normally reserved for the PLPA, read-only portions of the nucleus surrounding the 16MB line, and optionally the FLPA and MLPA. It is designed to protect those portions of virtual storage which are shared among address spaces, but are meant to contain read-only modules.

LOW ADDRESS PROTECTION

The low address protection mechanism protects against any program storing into addresses 0-511, where the processor and other hardware maintain status information. It can be turned on and off by manipulating a bit in one of the processor's control registers. Manipulation of this bit is usually reserved for functions like interrupt processing.

KEY-CONTROLLED PROTECTION

Storage keys have been a part of the 360/370 architecture since its inception several decades ago. They were initially the only protection mechanism available to shield one user's storage from another user's instructions. The multiple address space architecture introduced with MVS/370 largely supplanted the need for a protection key mechanism. However, it is still important for helping isolate various subsystems' code and data in areas like the CSA, where all users share a common addressing scheme, and, without protect keys, would have no means of assuring subsystem isolation.

Until recently, protect keys and virtual memory could be oriented toward either 2KB or 4KB page sizes. For each 2 or 4 KB block of memory, the operating system would assign a particular protect key. Subsequent attempts to access a block of memory would be permitted only by either the operating system or by a program operating with a protect key

that matched the memory being accessed.

With the advent of larger memory, the option to use 2KB pages and protect keys based upon them has been eliminated from the corresponding hardware. The 2KB protection option was eliminated early in the 308x series of processors.

The key-controlled protection mechanism works in this way: Each 4KB block of storage is assigned a 4-bit protection key and a fetch-protection bit. Any program will be assigned a 4-bit access key. If a program's access key is zero, it will bypass all key-controlled validity checks done by the hardware. If a program is attempting to fetch from a page whose fetch protection bit is off, the hardware will permit this, even if the program's access key is non-zero. If, however, a program whose access key is non-zero is attempting to store into the page, or the page's fetch protection bit is set on, then the program's access key must match the protection key of that page of memory. If it does not, the access will be prevented.

Figure 2-177 shows the protect keys normally used.

Storage	Storage	
Protect	Protect Key	
Key	Normally Used by	
0	MVS/XA System Control Program	
1	Job Scheduler/JES2/JES3	
2	VSPC	
3-4	Reserved	
5	Data Management	
6	TCAM/VTAM	
7	IMS	
8	V=V Problem Programs	
9-F	V=R Problem Programs	

Figure 2-177. Protect Keys

2.5.4.6 Virtual Storage Measurement by RMF

RMF does not actually monitor the virtual storage for which GETMAINS have been issued. Instead, it monitors via the VSMLIST macro the allocated areas, unallocated areas, and free areas of virtual storage. Thus, all RMF measurements of allocated and unallocated virtual storage will be in 4KB increments. Although the VSMLIST macro does report free areas within individual virtual pages which are not completely filled by the results of GETMAINS, RMF ignores this level of fragmentation and considers such a page to be completely allocated. Therefore, RMF monitors virtual storage at the same level of detail that the VSM manages it during its allocate/free processing of one or more 4KB pages.

Since VSM manages SQA by considering any unallocated SQA storage to be part of subpool 245 (see the discussion of SQA subpools above), RMF must measure SQA usage somewhat differently from the way it measures other areas of virtual storage. To measure other areas of virtual storage which are unallocated, RMF issues the VSMLIST macro with the SPACE=FREE parameter. Then it scans the results of this request, ignoring free space areas less than 4KB, eventually reporting the results as multiples of 4KB blocks of storage.

In contrast, when measuring areas other than SQA, RMF issues the VSMLIST macro with the SPACE=UNALLOC parameter to obtain a list of totally unallocated blocks of virtual memory. Each of these blocks represents one or more 4KB pages, so the results returned by the VSMLIST macro to RMF are already measured in multiples of 4KB.

Notice that RMF and VSM (and its VSMLIST reporting macro) use somewhat conflicting nomenclature. What RMF calls "free" storage is measured (except for SQA) by using the VSMLIST SPACE=UNALLOC parameter, not the SPACE=FREE parameter.

Also notice that RMF will generally understate the availability of virtual memory by some unknown amount, since what it calls allocated storage may in fact have holes of free memory contained in it. ("Free" as used here has the same meaning that VSM and the VSMLIST macro use.) Each such hole would be less than 4KB in size and would exist in a page that was already allocated by VSM to a particular subpool and protect key, so it could only be used to satisfy GETMAINS which matched these criteria.

Figure 2-178 summarizes the areas whose subpools are monitored by RMF.

Area	Subpools
SQA/ESQA	226, 239, 245
CSA/ECSA	227, 228, 231, 241
LSQA/ELSQA	255
PVT/EPVT	0, 1-127, 229, 230, 236, 237, 251, 252

Figure 2-178. Subpools Monitored by RMF

RMF monitors virtual storage in CSA and SQA at the level of protect keys and storage pools.

2.5.4.7 Allocation Measurements and Their Interpretation

For portions of virtual storage like CSA, SQA, and the User Region, a series of different measurements are available in the type 78, subtype 2, Monitor I Activity record written by RMF. This subtype record contains one Common Storage data section in which are recorded systemwide measures, and possibly one or more Private Area data sections if RMF has been instructed to track the use of virtual storage by particular jobs.

The following sections discuss the types of measurements recorded by RMF virtual storage monitoring.

SIZE OF XXX

Size of xx refers to the size limit in KB for xxx, where xxx represents CSA, SQA, User Region, or some other named portion of virtual storage for which MVS sets a size limit. This measurement represents the maximum size specified by the installation in SYS1.PARMLIB, via exits, or in some other fashion for this particular named portion of virtual storage.

ALLOCATED PAGES

Allocated Pages refers to the number of KB of storage allocated by the Virtual Storage Manager for GETMAINS within xxx, where xxx represents CSA, SQA, User Region, or some other named portion of virtual storage.

These allocations occur in 4KB increments. One or more GETMAINS may be satisfied by a single allocated block of virtual storage.

For a GETMAIN to succeed, there must be a contiguous block of virtual storage large enough to satisfy the GETMAIN. If there is currently no such block allocated, the Virtual Storage Manager (VSM) will attempt to find a free block of virtual storage that is large enough, and will allocate this block of storage.

If VSM is unable to find an allocated block with sufficient space which has not yet been GETMAINED or a free block of sufficient size, the function which depends upon this GETMAIN may fail. If the portion of virtual storage required is in a user address space, the user's job step may abnormally terminate. If the portion of virtual storage required is from a critical system area like CSA or SQA, then MVS may invoke emergency swapping until this shortage is relieved.

FREE PAGES

Free Pages refers to the number of KB of storage not currently allocated by the Virtual Storage Manager for GETMAINS within xxx, where xxx represents CSA, SQA, User Region, or some other named portion of virtual storage. These unallocated portions of virtual storage occur in 4KB increments. See the distinction described earlier for the meaning of this measurement for SQA.

ALLOCATED AREA

The Allocated Area represents the number of KB between the lowest allocated address of virtual storage and the highest allocated address within xxx, where xxx represents CSA, SQA, User Region, or some other named portion of virtual storage. This measure will include all free pages that lie between allocated pages.

This measurement will always be at least as large as the measure Allocated Pages, since it includes all allocated pages, plus any free pages scattered among those which are allocated. The degree to which the measure Allocated Area exceeds the measure Allocated Pages is an indication of fragmentation of virtual storage.

LARGEST FREE BLOCK

The Largest Free Block refers to the number of KB of virtual storage contained in the largest continuous block of storage not currently allocated by the Virtual Storage Manager for GETMAINS within xxx, where xxx represents CSA, SQA, User Region, or some other named portion of virtual storage.

Since any GETMAIN requires a contiguous virtual address range, this measure represents the size of the largest GETMAIN that could succeed within xxx.

A simplified example of these measurements, using a hypothetically very small CSA, is depicted in Figure 2-179, which shows the following measurements:

- o Size Of CSA 52 K bytes
- o Allocated Pages of CSA 28 K bytes
- o Free Pages of CSA 24 K bytes
- o Allocated Area of CSA 44 K bytes
- o Largest Free Block of CSA 12 K bytes




Figure 2-179. Allocation Measurements

MISCELLANEOUS MEASUREMENT CONSIDERATIONS

Most of the measurements described above are obtained via RMF interval driven sampling as minimums, maximums, and totals. Where minimums and maximums are maintained, there are corresponding time stamps to indicate when the minimum or maximum occurred.

RMF does not ordinarily sample virtual storage utilization as frequently as it samples other quantities in the system. The default RMF virtual storage sampling interval is set to be ten times as long as the RMF sampling interval for other measurements.

Totals can be used in conjunction with a data element containing the number of samples to calculate averages for the measurement period.

Since RMF samples its selected jobs for the Private Area data sections only when that job is executing and swapped in, various Private Area data sections may have different sample counts. The field SMF78SAM, found in the RMF Product section of the record, should be used when calculating averages from totals in the Common Storage data section. In contrast, each Private Area data section contains its own sample count, R782SAMP, used which should be used when calculating averages based upon totals in any particular Private Area data section.

There is one other concern relating to sampling intervals. Many of the values accumulated during the RMF sampling process for virtual storage monitoring are accumulated by RMF as single precision floating point numbers. The nature of floating point numbers dictates that if a relatively small number is added to a very large number, because of scale differences between the two numbers, some or all of the significance of the smaller number may be lost. This could lead to understating various sampled values if the RMF virtual storage sampling interval is set to too low a value.

Setting the sample interval to too low a value would cause the accumulated totals to become larger as more frequent samples are added to them, eventually leading to a substantial difference in magnitude between the total being accumulated and any given sample value being added to it. It is partly for this reason that the default RMF virtual storage sampling interval is set to be ten times as long as the RMF sampling interval for other measurements. Of course,

reducing the RMF virtual storage sampling interval will also increase the processor overhead due to sampling.

2.5.5 Case Study

The following case study illustrates the use of the Virtual Storage Management Analysis routines in evaluating the storage requirements of a system running IMS, TS0, and batch workloads. The software is used to review the last six months' activity in an effort to evaluate the impact of workload migration and recent upgrades in system software. The System Storage Usage and Common Area Usage Reports are used in the analysis.

The following topics are discussed:

- 1 - Problem Description
- 2 - Inquiry Panels
- 3 - Output Reports

2.5.5.1 Problem Description

Virtual storage resource usage should be monitored on an ongoing basis. Unlike CPU time or DASD space consumption, use of virtual storage is much more subtle. More and more storage is quietly consumed until some limit is reached with serious results, such as IMS Control region ABENDs or the system being quiesced and no new work being permitted to start. Also, depending on the allocation and subpool being used, expanding the use of virtual storage may also impact the use of real storage, as some subpools are defined as non-pageable. Finally, although some virtual storage constraint relief has been obtained through the introduction of the MVS/XA operating system, it is often necessary to modify applications and/or install new levels of software, such as IMS, to take advantage of the extended capabilities. Even though MVS/XA provides some virtual storage relief, it must be realized that not all applications can take advantage of the extended capabilities, and, therefore, virtual storage constraint will remain an issue.

The purpose of this study is to review the virtual storage usage of a system running IMS, TSO, and batch workloads. The site has a single 3090-400 system, operating in partition mode as two 3090-200 processors. Over the past several months the workloads have been shifted between the two systems in an effort to better balance the CPU requirements. We are using the VSM software to analyze the virtual storage requirements on the production system.

In particular, it must be determined if the Common Area allocations, such as SQA and CSA, are properly sized, what sort of growth patterns, if any, are present in the Common Area allocations, and what tasks are significant users of the Common Area storage.

2.5.5.2 Inquiry Panels

This section illustrates the inquiry panel entries we made to create the VSM System Storage and Common Area Storage Usage Reports to conduct our case study.

SYSTEM STORAGE USAGE REPORT

Figure 2-180 shows the inquiry panel that we used to create the System Storage Usage Report. (The resulting System Storage Usage Report is shown in Figure 2-182 in Section 2.5.5.3.) Section 2.5.6.1 presents a detailed description of each field in the inquiry panel.

The control panel shown in Figure 2-180 specifies that a System Storage Usage Report is to be created. This report is designed to assist you in determining the usage of virtual storage for the system under study. On this panel we selected the SYSID of the production system and also chose Zone 1, as this zone represents the prime shift for this site. Six cycles of the MONTHS timespan of the SCPVSM data were used in this analysis, as the changes outlined in the previous section all took place in this time period. Other zones could have been included in the study, but virtual storage issues affect the prime shift mainly due to the heavy usage of online systems.

COMMON AREA STORAGE REPORTS

Figure 2-181 show the inquiry panel that we used to create the Common Area Storage Usage Report. (The resulting Common Area Storage Usage Report is shown in Figure 2-183 in Section 2.5.5.3.) Section 2.5.6.2 presents a detailed description of each field in the inquiry panel.

The control panel shown in Figure 2-181 specifies that a Common Area Storage Usage Report is to be created. This report is designed to assist you in determining the usage of virtual storage in the Common Area for the system under study. This particular report was chosen because it amplifies the data shown on the System Storage Usage Report and permits additional analysis of the virtual storage usage issues. Because we wanted to parallel the reports, we used the same data selection criteria for this report as we did for the System Storage Usage Report. Six cycles of the MONTHS timespan of the SCPVSM data were also used in this analysis, as the changes outlined in the previous section all took place in this time period.

```

/----- VSM System Storage Usage Report -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: VMSYS - VSM Case Study - System Storage Usage |
| Timespan ==> MONTHS (DETAIL/DAYS/WEEKS/MONTHS)      Cycle(s) ==> 01 - 06 |
| Data base ==> U (MICSTEST) _                        _ |
| Allow run-time override of database ==> N (Y/N) and/or cycle(s) ==> N (Y/N) |
| Report by Hour if timespan is DAYS or WEEKS ==> N (Y/N) |
| Generate a report for each value of Zone ==> N (Y/N) |
| Report Minimum and Maximum values ==> Y (Y/N) |
| Data selection:                                     Allow run-time override |
|   Date ==> _ _ _ _ - _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> _ (Y/N) |
|   Hour ==> _ _ - _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> _ (Y/N) |
|   Zone ==> 1 _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> _ (Y/N) |
|   Sysid ==> SYSA _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> _ (Y/N) |
| Specify Extended Options ==> N (Y/N) |
\-----/

```

Figure 2-180. Case Study - System Storage Usage Report Inquiry Panel


```

/----- VSM Common Area Storage Usage Report -----/
| Command ==>
|
| Composing CA MICS Inquiry: VSMCSA - VSM Case Study - Common Area Storage Usage
|
| Timespan ==> MONTHS (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 06
| Data base ==> U (MICSTEST) _ _ _ _
| Allow run-time override of database ==> N (Y/N) and/or cycle(s) ==> N (Y/N)
|
| Report by Hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone ==> N (Y/N)
| Report Minimum and Maximum values ==> Y (Y/N)
| Report by Subpool ==> N (Y/N)
|
| Data selection:                                         Allow run-time override
|   Date ==> _ _ _ _ - _ _ _ _ - _ _ _ _ ==> _ (Y/N)
|   Hour ==> _ _ - _ _ - _ _ - _ _ ==> _ (Y/N)
|   Zone ==> 1 _ _ _ _ _ _ _ _ ==> _ (Y/N)
|   Sysid ==> SYSA _ _ _ _ _ _ _ _ ==> _ (Y/N)
|
| Specify Extended Options ==> N (Y/N)
|
/-----/

```

Figure 2-181. Case Study - Common Area Storage Usage Report Inquiry Panel

2.5.5.3 Output Reports

The execution of inquiries defined in the section above produced the excerpts of reports shown in Figures 2-182 and 2-234. The System Storage Usage Report will be discussed first, with the Common Area Usage Report being used to amplify the reported results where appropriate.

Several interesting observations can be made about the activity of the production system based on the two reports. First, the specified extended CSA storage allocation has nearly doubled in the reporting period. At the same time, the extended CSA storage actually allocated has dropped by 17%. The CSA size was increased in February to 12M and again in March to 16M, although the original specification in December would have been sufficient.

The second issue relates to the SQA EXPANSION TO CSA, found in the section of the report subtitled KEY AVG. VALUES - BELOW 16M. The SQA is expanding from 67K to 100K, depending upon the month. Referring to the report section KEY MIN/MAX VALUES - < 16M, the SQA actually expanded into CSA as much as 264K. Since this expansion seems to occur frequently, it might be advisable to expand the SQA specification. If the Private Area size below 16M is critical, then the CSA specification could be reduced. Care must be exercised here, as the CSA storage actually allocated is growing slowly.

In reviewing the Common Area Storage Report, several observations can be made. Although the extended CSA size was increased considerably, actual storage allocated was reduced. From the report section called CSA AVERAGE BY KEY - > 16M, we can see that Key 1 storage allocation has been reduced by over 98%. The Key 1 storage is used primarily by JES 2, and this system formerly had a large batch workload managed by RJE sites. Migrating this workload to another system and respecifying the JES parameters reduced CSA storage requirements considerably. At the same time, Key 6 storage allocations below 16M were reduced by over 70%, while increasing above 16M. There was no Key 6 storage allocation above 16M until February. This change coincided with the installation of a new VTAM release, which began using CSA storage above 16M. The largest increase in CSA storage usage occurred above 16M, in Key 7 allocation. This key is used by IMS, and this increase represented changes in the IMS configuration, as well as the migration of the production IMS workload to this system.

The extended CSA size is sufficient at this time. Indeed,

the extended CSA size could be reduced considerably with no ill effects if it were desirable to do so. Although overspecifying the CSA size causes no performance degradation, having excess CSA storage specified could lull one into a false sense of security and result in monitoring not being done often.

MONTHLY SYSTEM STORAGE USAGE REPORT							
SYSTEM = SYSA		RUNDATE = JUL 5, 1988					
MONTHS (YY/MM):	87-12	88-01	88-02	88-03	88-04	88-05	CHANGE %
AVERAGE MAP - ABOVE 16 MEG							
PRIVATE	2009M	2009M	2004M	2001M	2001M	2001M	-0.42 %
CSA	7998K	8548K	12M	16M	16M	16M	106.12 %
MLPA	44K	44K	44K	44K	44K	44K	---
FLPA	16K	16K	16K	16K	16K	16K	---
PLPA	3968K	3968K	4096K	4220K	4220K	4220K	6.35 %
SQA	8512K	8512K	8512K	8512K	8512K	8512K	---
NUCLEUS	2464K	2464K	2464K	2464K	2464K	2464K	---
AVERAGE MAP - BELOW 16 MEG							
NUCLEUS	412K	412K	409K	418K	424K	424K	2.91 %
SQA	640K	640K	640K	640K	640K	640K	---
PLPA	2541K	2520K	2406K	2306K	2312K	2312K	-9.01 %
FLPA	4K	4K	4K	4K	4K	4K	---
MLPA	4K	5K	4K	4K	4K	4K	---
CSA	6638K	6658K	6775K	6866K	6856K	6856K	3.26 %
PRIVATE	6140K	6140K	6140K	6140K	6140K	6140K	---
PREFIX SAVE AREA	4K	4K	4K	4K	4K	4K	---
CSA AVG VALUES - ABOVE 16M							
FREE PAGES (BYTES)	3776K	3643K	9864K	13M	13M	13M	268.19 %
LARGEST FREE BLOCK	3684K	3441K	9524K	12M	12M	12M	251.03 %
ALLOCATED AREA SIZE	4319K	5219K	3057K	3996K	3507K	3554K	-17.70 %
CSA AVG VALUES - BELOW 16M							
FREE PAGES (BYTES)	1222K	1131K	1549K	1816K	1739K	1674K	36.94 %
LARGEST FREE BLOCK	1065K	966K	1406K	1691K	1572K	1536K	44.24 %
ALLOCATED AREA SIZE	5646K	5715K	5712K	5171K	5280K	5809K	2.89 %

Figure 2-182. Case Study - System Storage Usage Report (Part 1 of 3)

MONTHLY SYSTEM STORAGE USAGE REPORT							
SYSTEM = SYSA		RUNDATE = JUL 5, 1988					
MONTHS (YY/MM):	87-12	88-01	88-02	88-03	88-04	88-05	CHANGE %
SQA AVG VALUES - ABOVE 16M							
FREE PAGES (BYTES)	6476K	6495K	6578K	6518K	6530K	6590K	1.74 %
LARGEST FREE BLOCK	6405K	6447K	6516K	6457K	6462K	6545K	2.19 %
ALLOCATED AREA SIZE	8512K	8512K	8512K	8512K	8512K	8512K	---
SQA AVG VALUES - BELOW 16M							
FREE PAGES (BYTES)	13K	8K	11K	6K	9K	9K	-31.21 %
LARGEST FREE BLOCK	5K	4K	4K	4K	4K	4K	-26.98 %
ALLOCATED AREA SIZE	638K	638K	638K	638K	638K	638K	-0.10 %
KEY AVG VALUES - ABOVE 16M							
SQA EXPANSION TO CSA	0K	0K	0K	0K	0K	0K	N/A
PLPA PACKING LOSS	9K	9K	10K	11K	11K	10K	15.49 %
PLPA REDUNDANT SPACE	13K	13K	13K	13K	13K	13K	---
MAXIMUM USER REGION	2000M	2000M	1996M	1992M	1992M	1992M	-0.42 %
KEY AVG VALUES - BELOW 16M							
SQA EXPANSION TO CSA	67K	83K	76K	100K	95K	88K	31.05 %
PLPA PACKING LOSS	6K	5K	6K	6K	6K	6K	-8.68 %
PLPA REDUNDANT SPACE	3K	4K	2K	2K	2K	2K	-12.82 %
MAXIMUM USER REGION	5988K	5988K	5988K	5988K	5988K	5988K	---

Figure 2-182. Case Study - System Storage Usage Report (Part 2 of 3)

MONTHLY SYSTEM STORAGE USAGE REPORT							
SYSTEM = SYSA		RUNDATE = JUL 5, 1988					
MONTHS (YY/MM):	87-12	88-01	88-02	88-03	88-04	88-05	CHANGE %
SQA MIN/MAX VALUES - < 16M							
FREE PAGES(BYTES) MIN	0K	0K	0K	0K	0K	0K	N/A
FREE PAGES(BYTES) MAX	204K	232K	228K	184K	192K	224K	9.80 %
LARGEST FREE BLK MIN	0K	0K	0K	0K	0K	0K	N/A
LARGEST FREE BLK MAX	204K	200K	180K	164K	192K	176K	-13.72 %
ALLOC. AREA SIZE MIN	436K	440K	460K	476K	448K	464K	6.42 %
ALLOC. AREA SIZE MAX	640K	640K	640K	640K	640K	640K	---
KEY MIN/MAX VALUES - > 16M							
SQA EXP. INTO CSA MIN	0K	0K	0K	0K	0K	0K	N/A
SQA EXP. INTO CSA MAX	0K	0K	0K	0K	0K	0K	N/A
PLPA PACKING LOSS MIN	9K	9K	9K	11K	11K	10K	14.74 %
PLPA PACKING LOSS MAX	9K	9K	11K	11K	11K	11K	18.26 %
REDUNDANT PLPA MIN	13K	13K	13K	13K	13K	13K	---
REDUNDANT PLPA MAX	13K	13K	13K	13K	13K	13K	---
MAX. USER REGION MIN	2000M	2000M	1992M	1992M	1992M	1992M	-0.39 %
MAX. USER REGION MAX	2005M	2000M	2000M	1992M	1992M	1992M	-0.64 %
KEY MIN/MAX VALUES - < 16M							
SQA EXP. INTO CSA MIN	0K	0K	0K	0K	0K	0K	N/A
SQA EXP. INTO CSA MAX	156K	264K	188K	200K	172K	172K	10.25 %
PLPA PACKING LOSS MIN	5K	5K	5K	6K	6K	6K	3.50 %
PLPA PACKING LOSS MAX	8K	5K	8K	7K	6K	6K	-29.30 %
REDUNDANT PLPA MIN	2K	2K	2K	2K	2K	2K	---
REDUNDANT PLPA MAX	4K	39K	2K	2K	2K	2K	-37.31 %
MAX. USER REGION MIN	5988K	5988K	5988K	5988K	5988K	5988K	---
MAX. USER REGION MAX	5988K	5988K	5988K	5988K	5988K	5988K	---

Figure 2-182. Case Study - System Storage Usage Report (Part 3 of 3)

MONTHLY COMMON AREA STORAGE USAGE REPORT BY KEY							
SYSTEM = SYSA RUNDATE = JUL 5, 1988							
MONTHS (YY/MM):	87-12	88-01	88-02	88-03	88-04	88-05	CHANGE %
CSA AVERAGE BY KEY - > 16M							
KEY 0 ALLOCATION	527K	570K	541K	593K	603K	583K	10.57 %
KEY 1 ALLOCATION	2916K	3431K	953K	32K	32K	32K	-98.90 %
KEY 2 ALLOCATION	0K	0K	0K	0K	0K	0K	N/A
KEY 3 ALLOCATION	57K	59K	57K	64K	63K	63K	11.45 %
KEY 4 ALLOCATION	0K	0K	0K	0K	0K	0K	N/A
KEY 5 ALLOCATION	57K	56K	57K	56K	55K	55K	-2.15 %
KEY 6 ALLOCATION	0K	0K	338K	649K	616K	608K	N/A
KEY 7 ALLOCATION	608K	709K	706K	927K	971K	1183K	94.38 %
KEYS 8-F ALLOCATION	55K	77K	62K	93K	69K	56K	1.86 %
CSA AVERAGE BY KEY - < 16M							
KEY 0 ALLOCATION	1619K	1626K	1629K	1644K	1632K	1615K	-0.25 %
KEY 1 ALLOCATION	171K	153K	150K	158K	154K	153K	-10.72 %
KEY 2 ALLOCATION	0K	0K	0K	0K	0K	0K	N/A
KEY 3 ALLOCATION	99K	116K	94K	110K	109K	117K	17.99 %
KEY 4 ALLOCATION	0K	0K	0K	0K	0K	0K	N/A
KEY 5 ALLOCATION	22K	24K	23K	24K	21K	21K	-4.34 %
KEY 6 ALLOCATION	622K	677K	409K	180K	168K	167K	-73.12 %
KEY 7 ALLOCATION	2683K	2715K	2719K	2702K	2813K	2895K	7.90 %
KEYS 8-F ALLOCATION	130K	129K	122K	128K	121K	123K	-5.30 %

Figure 2-183. Case Study - Common Area Storage Usage Report

2.5.6 Component Operation

This component operation section describes the operation of the Virtual Storage Analysis routines in terms of the CA MICS Performance Management Option inquiry options selected.

Inquiry options available with the Virtual Storage Analysis software describe the reports to be generated. In particular, these options define:

- o The cycles from the CA MICS RMF files to be processed.
- o The summarization level of the reports as hourly, daily, weekly, or monthly. The availability of each summarization level assumes that your system supports the required files at each level.
- o The reports and plots to be produced.

The following fields are common to most control panels:

Timespan:	The timespan (DETAIL, DAYS, WEEKS, or MONTHS) in the CA MICS file(s) that you wish to use as input for processing.
Cycle(s):	One or more cycles of the CA MICS files that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a CYCLE from the CA MICS file will usually contain a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided by the DATE and HOUR control options.
Database:	The single-character CA MICS identifier for the database to be accessed. Valid values are the database unit identifiers available at your installation. Contact your CA MICS System Administrator if you do not know what database units are available, or enter a question mark (?) to view a list of your database units. You can select up to five databases per inquiry.
Allow Run-Time Override:	The indication of whether the database identifier and/or the Cycles specification can be overridden at execution time.
Report by	The indication of whether a report or plot

hour:	should be generated for each hour.
Generate a report for each value of Zone:	The indication of whether a separate report or plot should be generated for each zone defined in the input data or whether all zones should be merged together for the report or plot.
Date:	Up to two date ranges (DDMMYY format) to be used for analysis. This option is normally used to drop the extra day from a cycle that is due to the timing of the SMF dump.
Hour:	One to three ranges of hours for which data is to be reported. For an individual hour, leave the end hour blank or make it the same as the start hour. You can specify a value from 00 to 23.
Zone:	The specification that zone selection is to be applied. If zone 1 definition is 8:00 a.m. to 6:00 p.m., you can specify reports covering that time period by entering a 1 in the ZONE field.
SYSID:	The specification that SYSIDs are to be selected for reporting. You can enter up to six SYSIDs to be selected. If you leave the field blank, no testing of SYSIDs will be done.
Job Name:	The specific JOBNAME(s) to be reported. If you code this field, only those specified will be selected.
Program:	The specific PROGRAM(s) to be reported. If you code this field, only those specified will be selected.
Specify extended options:	The specification that allows you to supply additional report specifications. These options include special SAS data selection statements, non-standard input and output files, and special graphics options.

The operation of the Timespan, Cycle(s), and Zone fields depends on definitions made during the installation of CA MICS.

Figure 2-184 shows the Virtual Storage Management (VSM) Analysis menu that you use to select the desired report or plot within the VSM reporting area.

2.5.6.1 VSM System Storage Usage Report

The VSM System Storage Usage Report details the use of virtual storage by the entire system. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the VSM System Storage Usage Report obtains its data from the SCPVSM Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The VSM System Storage Usage Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options that are normally defaulted.

Figure 2-185 shows a VSM System Storage Usage Report inquiry panel. The fields common to all Virtual Storage Management Analysis routines are described in Section 2.5.6. The following are general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Report by hour and Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized to the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

There are four other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.5.6.2 VSM Common Area Storage Usage Report

The VSM Common Area Storage Usage Report details the use of virtual storage in the system Common Area. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the VSM Common Area Storage Usage Report obtains its data from the SCPVSM Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The VSM Common Area Storage Usage Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options that are normally defaulted.

Figure 2-186 shows a VSM Virtual Storage Common Area Usage Report inquiry panel. The fields common to all Virtual Storage Management Analysis routines are described in Section 2.5.6. The following are general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Report by hour and Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized to the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

There are four other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

```

/----- VSM Common Area Storage Usage Report -----/
|
| Command ==>
|
| Composing CA MICS Inquiry:  SYS001 - Common Area Storage Usage Report Inquiry
|
| Timespan ==> WEEKS  (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 06
| Data base ==> U (MICSTEST) _ _ _ _ _
| Allow run-time override of database ==> N (Y/N) and/or cycle(s) ==> N (Y/N)
|
| Report by Hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone      ==> N (Y/N)
| Report Minimum and Maximum values              ==> N (Y/N)
|
| Data selection:                                         Allow run-time override
|   Date          ==> _ _ _ _ - _ _ _ _ - _ _ _ _      ==> _ (Y/N)
|   Hour          ==> _ _ - _ _ - _ _ - _ _ - _ _      ==> _ (Y/N)
|   Zone          ==> _ _ - _ _ - _ _ - _ _              ==> _ (Y/N)
|   Sysid         ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _      ==> _ (Y/N)
|
|
| Specify Extended Options          ==> Y (Y/N)
|
/-----/

```

Figure 2-186. VSM Common Area Storage Usage Report Inquiry Panel

2.5.6.3 VSM Private Storage Usage Report

The VSM Private Storage Usage Report details the use of virtual storage by individual jobs. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the VSM Private Storage Usage Report obtains its data from the SCPVSU File. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The VSM Private Storage Usage Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options that are normally defaulted.

Figure 2-187 shows an VSM Private Storage System Usage Report inquiry panel. The fields common to all VSM Analysis routines are described in Section 2.5.6. The following are general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization to the ZONE level can be done for any timespan. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Report by hour and the Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized into the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, etc.).

If you choose the TRENDING report type, then you must specify a Trending Element. You can use the name of any numeric element present in the SCPVSU File. Enter a question mark (?) in this field to obtain a list of all the elements in the selected file.

Note that not all data selection values are valid for every type of input file, as some of the elements being tested do not exist in certain files. An error message is displayed on the panel if you specify an invalid selection.

There are four other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

```

/----- VSM Private Storage Usage Report -----/
| Command ==>
|
| Composing CA MICS Inquiry: PVT001 - Private Storage Usage Report Inquiry
|
| Timespan ==> WEEKS (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 06
| Data base ==> U (MICSTEST) _
| Allow run-time override of database ==> N (Y/N) and/or cycle(s) ==> N (Y/N)
|
| Report by Hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone ==> N (Y/N)
| Report Minimum and Maximum values ==> N (Y/N)
|
| Report Type ==> OVERVIEW (OVERVIEW/TRENDING)          Trending Element ==> _____
|
| Data selection:                                         Allow run-time override
|   Date ==> _____ - _____ - _____ ==> _ (Y/N)
|   Hour ==> __ - __ - __ - __ - __ - __ ==> _ (Y/N)
|   Zone ==> _ - _ - _ - _ - _ ==> _ (Y/N)
|   Sysid ==> _____ ==> _ (Y/N)
|   Job name ==> _____
|   Program name ==> _____
|
| Specify Extended Options ==> N (Y/N)
/-----/

```

Figure 2-187. VSM Private Storage Usage Report Inquiry Panel

2.5.6.4 VSM System Storage Usage Plot

The VSM System Storage Usage Plot shows the use of virtual storage by the entire system. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the VSM System Storage Usage Plot obtains its data from the SCPVSM Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The VSM System Storage Usage Plot is produced using a series of six panels. Only the first panel is required. The others are used to specify additional options that are normally defaulted.

Figure 2-188 shows a VSM Virtual Storage System Usage Plot inquiry panel. The fields common to all Virtual Storage Management Analysis routines are described in Section 2.5.6. The following are general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Plot by hour and Generate a plot for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate plot is produced for each day. If you do not specify a summarization parameter, the plot is summarized to the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

There are five other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations
- o Extended Color Options Menu

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

```

/----- VSM System Storage Usage Plot -----\
| Command ==>                                     |
| Composing CA MICS Inquiry:  SYS003 - System Storage Usage Plot Inquiry |
| Timespan ==> WEEKS (DETAIL/DAYS/WEEKS/MONTHS)      Cycle(s) ==> 01 - 06 |
| Data base ==> U (MICSTEST) _                      _ |
| Allow run-time override of database ==> N (Y/N) and/or cycle(s) ==> N (Y/N) |
| Only Plot the Below 16 meg. Storage ==> N (Y/N) |
| Generate color graphic plots          ==> N (Y/N) |
| Data selection:                                Allow run-time override |
|   Date          ==> _ _ _ _ - _ _ _ _ - _ _ _ _ ==> _ (Y/N) |
|   Hour          ==> _ _ - _ _ - _ _ - _ _ - _ _ ==> _ (Y/N) |
|   Zone          ==> _ _ _ _ _ _ _ _ _ _ ==> _ (Y/N) |
|   Sysid         ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> _ (Y/N) |
| Specify Extended Options          ==> N (Y/N) |
\-----/

```

Figure 2-188. VSM System Storage Usage Plot Inquiry Panel

2.5.6.5 VSM Private Storage Usage Plot

The VSM Private Storage Usage Plot details the use of virtual storage by individual jobs. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the VSM Private Storage Usage Plot obtains its data from the SCPVSU File. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The VSM Private Storage Usage Plot is produced using a series of six panels. Only the first panel is required. The others are used to specify additional options that are normally defaulted.

Figure 2-189 shows a VSM Private Storage System Usage Plot inquiry panel. The fields common to all VSM Analysis routines are described in Section 2.5.6. The following are general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization to the ZONE level can be done for any timespan. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Plot by hour and the Generate a plot for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate plot is produced for each day. If you do not specify a summarization parameter, the plot is summarized into the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, etc.).

The optional TRENDING element for plotting has no default value. You can use the name of any numeric element present on the SCPVSU File. Enter a question mark (?) in this field to obtain a list of all the elements on the selected file.

Note that not all data selection values are valid for every type of input file, as some of the elements being tested do not exist in certain files. An error message is displayed on the panel if you specify an invalid selection.

There are five other panels in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu

- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations
- o Extended Color Graphic Options Menu

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

```

/----- VSM Private Storage Usage Plot -----\
| Command ==>
|
| Composing CA MICS Inquiry:  PVT002 - Private Storage Usage Plot Inquiry
|
| Timespan ==> WEEKS (DETAIL/DAYS/WEEKS/MONTHS)      Cycle(s) ==> 01 - 06
| Data base ==> U (MICSTEST) _
| Allow run-time override of database ==> N (Y/N) and/or cycle(s) ==> N (Y/N)
|
| Plot Type ==> OVERVIEW (OVERVIEW/TRENDING)      Trending Element ==> _____
|
| Generate color graphic plots ==> N (Y/N)
|
| Data selection:                                     Allow run-time override
|   Date           ==> ____ - ____ - ____          ==> _ (Y/N)
|   Hour           ==> __ - __ - __ - __ - __        ==> _ (Y/N)
|   Zone           ==> _ - _ - _ - _ - _             ==> _ (Y/N)
|   Sysid          ==> ____ - ____ - ____ - ____      ==> _ (Y/N)
|   Job name       ==> _____
|   Program name  ==> _____
|
| Specify Extended Options      ==> N (Y/N)
|
\-----/

```

Figure 2-189. VSM Private Storage Usage Plot Inquiry Panel

2.5.7 Requirements

The Virtual Storage Management Analysis reports and plots enable you to track the allocation of virtual storage on a monthly, weekly, daily, hourly, or interval basis. Virtual storage requirements at the job or step level are reported to help you in evaluating the impact of new applications or users, or the growth of existing applications.

Figure 2-190 shows the CA MICS files and timespan levels that must be maintained at your installation in order to produce all of the Virtual Storage Management Analysis reports and plots.

CA MICS FILE	LEVEL				
	X	D	W	M	Y
SCPVSM	x	x	x	x	
SCPVSU	x	x	x	x	

Figure 2-190. VSM Analysis CA MICS File Requirements

The following is a list of files required to run each individual report:

- o System Storage Usage Report - SCPVSM
- o Common Area Storage Usage Report - SCPVSM
- o Private Storage Usage Report - SCPVSU
- o System Storage Usage Plot - SCPVSM
- o Private Storage Usage Plot - SCPVSU

Only one level is required for the CA MICS RMF product files (i.e., DETAIL, DAYS, WEEKS, or MONTHS), depending on the timespan you specify during inquiry definition.

The number of cycles that are maintained at your installation for each of the CA MICS files in Figure 2-190 determines the time frame for which you can produce VSM Analysis reports and plots.

2.6 PR/SM LPAR Performance Analysis

This section introduces the concepts underlying IBM's Processor Resource/Systems Manager (PR/SM) feature as background for discussing performance management in this environment. The composition of a Central Electronic Complex (CEC) is described, emphasizing those hardware and software components to monitor and adjust to maximize throughput and minimize response time for online systems.

PR/SM allows a CEC to be divided up into one or more logical partitions (LPAR); each partition is independent of the others. PR/SM allows multiple operating System Control Programs (SCPs), such as MVS/XA, MVS/ESA, VM/SP, or VM/XA to share the facilities of a single CEC or one side of a physically partitioned complex. A CEC or physical partition can be initialized either in LPAR mode or in basic mode. In basic mode, the functions of PR/SM are not available, so the following discussion addresses LPAR mode only.

The PR/SM Performance Analysis inquiries let you track the use of CECs and the activity of individual LPARs at any supported timespan and summarization pattern. Reports display system measurements on overall CEC utilization and configuration as well as the activity of individual LPARs and their proportionate share of a complex.

Section 2.6.2, PR/SM LPAR Concepts, defines the terms and explains the basic dynamics of many of the elements used in PR/SM performance analysis. If you are familiar with with these terms and concepts, you must proceed directly to Section 2.6.3.

The following sections describe the operation of the PR/SM LPAR Analysis inquiries and how to use the reports that are produced:

- 1 - Functional Description
- 2 - PR/SM LPAR Concepts
- 3 - Standard Output
- 4 - Usage Guidelines
- 5 - Reserved
- 6 - PR/SM LPAR Analysis Operation
- 7 - Requirements

2.6.1 Functional Description

The PR/SM LPAR Analysis inquiries allow you to track the utilization and configuration of CECs and the LPARs into which each complex is subdivided. Any supported CA MICS timespan and summarization pattern can be implemented.

Explanations of the terms used below in describing the reports produced by these functions are presented in Section 2.6.4. A description of the controls available and how to evaluate the effect of those controls with the PR/SM Analysis reports is given in the standard output and PR/SM LPAR concepts sections (see Sections 2.6.3 and 2.6.4).

Two available PR/SM reports are:

- o CEC Summary Report - This report can be produced for each user-selected CEC, tracking utilization and environmental measures. The percent of the CEC consumed by all LPARs is reported to provide an indication of whether additional LPARs can be absorbed by the complex and to aid in determining if the CEC is configured to achieve optimum throughput. To aid in such analysis, CEC configuration information is also reported.
- o LPAR Activity Report - This report can be produced for each user selected logical partition, tracking utilization and environmental measures. Utilization is shown in terms of the percentage of the total CEC consumed by the LPAR, and the percentage of shared system and logical processor utilization for the LPAR. This allows for analysis of what load a given LPAR is placing on the entire CEC. These measures enable you to analyze the load an LPAR is placing on the CEC. They also indicate whether additional processor resource should be allocated to a given LPAR, or whether they should be given up by this LPAR for use by other LPARs. To aid in such analysis, configuration data is also reported.

Each report can be tailored during online definition to display data at various timespans and summarization patterns (reporting by zone, hour, or day, etc.). Overall CEC Summary analysis, as well as the activity of individual LPARs can be tracked via these report.

These reports utilize information derived from the CA MICS HARLPC file. The range of data that can be reported from a

CA MICS database depends on the timespans that were defined for the HARLPC file and the number of cycles defined in each timespan.

2.6.2 PR/SM LPAR Concepts

The PR/SM feature is a hardware facility that was made available by IBM starting with the IBM 3090E and 3090S series processors. It is an intermediate step between physical (or hardware) partitioning, such as the reconfiguration of a 3090-400E into the equivalent of two separately IPLable 3090-200E systems, and software partitioning, as implemented in the Virtual Machine Facility (VM).

PR/SM allows multiple operating System Control Programs (SCPs), such as MVS/XA or MVS/ESA, to share the facilities of a single Central Electronic Complex (CEC) or one side of a physically partitioned complex. A CEC or physical partition can be initialized either in logical partition (LPAR) mode or in basic mode. In basic mode, the functions of PR/SM are not available.

PR/SM hardware is utilized by two IBM products: a microcode feature called Logical Partitioning Mode and VM/XA Multiple Preferred Guest. The hardware feature can support either of these two products, but not both simultaneously. The discussion which follows addresses only logical partitioning mode operations.

The processor, storage, and I/O resources of the processor complex can be divided and allocated to several logical partitions (LPARs), which can then be IPLed and operated independently of each other. Up to ten LPARs per physical partition are allowed, meaning that a CEC that supports physical partitioning can have up to ten per side for a maximum of twenty LPARs, while a non-partitionable complex is limited to ten. Throughout this chapter, the terms CEC and complex are used interchangeably, and can be taken to refer to one half of a physically partitioned complex.

The following sections discuss PR/SM operations and controls allowing an analyst to optimally configure and operate a CEC:

- 1 - Overview of PR/SM LPAR Operation
- 2 - Controlling the PR/SM LPAR Environment

2.6.2.1 Overview of PR/SM LPAR Operation

When an LPAR is defined, the amount of central and expanded storage allocated to it is specified in one-megabyte increments. This storage is dedicated (not shared with nor accessible in any way by other LPARs).

Individual central processors (CPs), on the other hand, can either be dedicated to or shared among LPARs. Dedicated CPs are used exclusively by the LPAR and are not controlled in any way by the PR/SM feature. An LPAR can have multiple logical CPs defined, but the number cannot exceed the number of CPs under PR/SM control. In other words, the total number of CPs in the complex minus the number of dedicated processors, if any, is the number of multiple logical processors that can be defined to an LPAR.

PR/SM contains facilities to prioritize LPAR access to CPs during periods of contention. Within a single complex it is possible to mix LPARs, including dedicated processors with LPARs sharing processors, assuming sufficient CPs exist. For any LPAR with more than one defined logical processor, however, all of them must be dedicated or all must be shared.

Vector processing elements are associated with the CPs to which they are attached within the CEC, and can logically be thought of as being a part of the CP itself for purposes of this discussion.

The PR/SM feature allocates processor resources by assigning CPs to logical processors as needed by the active LPARs. If the total demand of active LPARs is less than the total resources of the complex, there is no limit placed on the availability of CPs to any LPAR other than the limit implied by the number of logical processors defined for the LPAR: there can never be more CPs allocated than the number of defined logical processors. When demand exceeds supply, PR/SM prioritizes CP allocations based on values specified when the LPAR was defined (this weighting can be modified dynamically by an authorized operator). Depending on your processor type, you may be able to cap LPARs to limit the percentage of the total processor it uses.

CPs are allocated to an LPAR by means of time slicing. The duration of the time slice can be determined dynamically by the PR/SM feature, and is based on the total number of non-dedicated CPs and the number of active logical processors. The duration of a time slice can range from a minimum of 12.5 milliseconds to a theoretical maximum of 300

milliseconds. A CP will remain allocated to an LPAR until either the time slice expires or the SCP running in the LPAR attempts to place the logical processor in the wait state; at this time, the CP will be made available for work elsewhere.

2.6.2.2 Controlling the PR/SM LPAR Environment

As was outlined in the previous section, PR/SM shares processor resources among LPARs via weighted prioritization and by time-slicing. Both of these controls can, however, be overridden by the console operator or system programmer. The weighting for each LPAR can be dynamically modified, and the controls over time-slicing can be altered by means of enabling the wait assist feature and setting a fixed time slice. When this is done, PR/SM will not take control of a CP when the logical processor goes into a wait state, but leaves the CP under the control of the LPAR for the entire duration of the time slice.

This level of control exacts a price: the PR/SM feature itself uses some of the processor resources that it is managing and allocating on behalf of the logical partitions. This overhead is usually quite small and is directly related to the ratio of active logical processors to the number of physical processors under PR/SM control. The use of dedicated processors greatly reduces this overhead, perhaps to as low as 0.1 percent. The overhead can grow to as much as 12 percent in some environments as more processors are shared and as contention for their use increases.

LPARs with dedicated processors have such low overhead that in some circumstances it is possible to increase rather than decrease the throughput of a complex by using PR/SM. Since the overhead involved in managing, coordinating, and synchronizing multiple processors by a single SCP can be substantial, reducing or eliminating multiple CPs can result in substantial savings of processor resources. If several workloads currently supported simultaneously within a single multiple-processor system can be divided into multiple single-processor LPARs with dedicated CPs, up to a 10 percent increase in overall throughput is possible.

2.6.3 Standard Output

The PR/SM LPAR analysis reports give you the capability to track the utilization of CECs and LPARs at any supported time-span and summarization level. Both CECs and LPARs can be reported on from the perspective of different systems' (SYSIDs') recording of RMF data for the same complex.

Configuration information in terms of number of shared and dedicated processors and Wait State Assist status are reported to aid you in evaluating the impact of changes to the configuration, such as adding or removing LPARs, or enabling or disabling the Wait State Assist feature. You can use the selection parameters described in Section 2.6.6, Component Operation, to limit reporting to specific CECs, LPARs, systems, and times.

Described below are the report formats of the CEC Summary Report, and the LPAR Activity Report. Common title and report elements are also defined.

- 1 - Common Report Elements
- 2 - CEC Summary Report
- 3 - LPAR Activity Report

2.6.3.1 Common Report Elements

This section describes elements common to all PR/SM Analysis reports. Options that are unique to a specific report are documented with that report.

TITLE ELEMENTS

The following elements are in the standard title portion of each report:

Installation Name: The installation name that is defined at CA MICS system installation time. This value is centered in the first line of the report. For more information about this value, consult your CA MICS System Administrator.

Product Name: CA MICS Performance Manager

Report Name: The name of the report, centered in the third line of the heading. It cannot be modified.

SYSID or System: The CA MICS logical SYSID being reported. Data from multiple systems cannot be combined using standard control parameters.

LPAR Name: The name of the LPAR being reported. For the CEC Summary report, only one LPAR need be specified to produce the report. For the LPAR Activity Report each selected LPAR is reported, or all LPARs for a given SYSID.

ZONE: If data is being reported by zone, the zone number is shown.

RUNDATE: The date the report was run.

ADDITIONAL ELEMENTS

Reporting Interval: The first row of some reports shows the date or date/time combination of the period being reported. If the date being reported on is from the WEEKS timespan, the date reflects the starting date of the week, as derived from the CA MICS variable STARTTS. In all other cases, the date reflects the ending date, derived from the CA MICS variable ENDTS. In monthly

reports, only the year and month portions of the date are displayed.

% CHANGE: The last column on some reports, this field is the percent change for each report item (row) between the base period (first data column) and the last period (last data column). If data for only one period is reported, then dashes (---) are displayed. If the base period is zero, then N/A is printed to indicate that a percent change could not be calculated.

2.6.3.2 CEC Summary Report

The CEC Summary Report tracks the utilization and environmental configuration information at a monthly, weekly, daily, or hourly level. Figure 2-191 shows an example of a monthly CEC Summary Report and Figure 2-192 shows the hourly form of the report. In the hourly form, data for one day is displayed by hour using the DAYS timespan of the PR/SM LPAR Configuration and Activity File (HARLPC).

Because multiple LPARs can be collecting data (via RMF) on overall CEC utilization, it is possible to report on the same CEC activity derived from different SYSID/LPAR combinations. Generally speaking, CEC Summary reporting should utilize the SYSID/LPAR data that is most comprehensive (that is, the LPAR that was most active during the interval being reported).

REPORT OPTIONS

The CEC Summary Report supports standard time, SYSID, LPAR, and external data selection. These selections are controlled by the options discussed in Section 2.6.6.

DATA ELEMENTS

The elements reported are divided into the following categories:

- o Utilization Measures
- o Environment

UTILIZATION MEASURES

Utilization is shown in terms of the percentage of the total CEC consumed by all work running on that CEC. Items reported are:

CPU TM SHARD OR DED: The total CPU time used by both dedicated and shared processors on this CEC. For dedicated processors, this includes any wait time that the processor may have had since this time is not available to the rest of the CEC. It is derived by the following computation:

$$\text{PRSM TDTM} + ((\text{PRSM TDP} / \text{INTERVALS}) * \text{DURATION})$$

CPU TIME AVAILABLE: The total available CPU time across the CEC, including both dedicated and shared processors

as derived by the following computation:

$$(\text{PRSMTPP} / \text{INTERVALS}) * \text{DURATION}$$

PERCENT CEC CONSUMED: The percent of this CEC consumed by all LPARs. This includes the wait time of any dedicated processors since, from the viewpoint of the CEC as a whole, this is processor time that is not available (that is, is consumed). It is computed as the ratio of the two preceding report items times 100.

ZAAP TIME SHARED OR DEDICATED: The total CPU time used by both dedicated and shared Application Assist Processors (zAAPs) on this CEC. It is derived by the following computation.

$$\text{PRSMTZST} + ((\text{PRSMTZDP} / \text{RMFZAPAV}) * \text{DURATION})$$

ZAAP TIME AVAILABLE: The total available zAAP CPU time across the CEC including both dedicated and shared zAAP processors derived by the following computation:

$$(\text{PRSMTZP} / \text{RMFZAPAV}) * \text{DURATION}$$

PERCENT ZAAP PROCESSORS CONSUMED: The percent of this CEC consumed by all LPARs. This is processor time that is not available (that is, is consumed). It is computed as the ratio of the two preceding report items times 100.

ZIIP TIME SHARED OR DEDICATED: The total CPU time used by both dedicated and shared Integrated Information Processors (zIIP) on this CEC. It is derived by the following computation.

$$\text{PRSMTSST} + ((\text{PRSMTSDP} / \text{RMFSUPAV}) * \text{DURATION})$$

ZIIP TIME AVAILABLE: The total available zIIP CPU time across the CEC including both dedicated and shared zIIP processors derived by the following computation:

$$(\text{PRSMTSP} / \text{RMFSUPAV}) * \text{DURATION}$$

PERCENT ZIIP PROCESSORS The percent of this CEC consumed by all LPARs. This is processor time that is not available (that is, is consumed). It is computed as the

CONSUMED: ratio of the two preceding report items times 100.

ENVIRONMENT

The environmental measures appearing in this report reflect the CEC configuration and monitoring activity.

OF PHYS PROCESSORS: The number of physical processors assigned for use by PR/SM on this CEC: PRSMNOPP.

AV # OF DED LOG PROCS: The average number of processors dedicated to one or more LPARs on the CEC as derived in the following computation:

$$\text{PRSM TDP} / \text{PRSM USCT}$$

AV # OF DEFINED LOG PARTS: The average number of defined logical partitions on the CEC as derived in the following computation:

$$\text{PRSM T LPS} / \text{PRSM USCT}$$

AV # OF ACTIVE PARTS: The average number of active logical partitions in the CEC as derived in the following computation:

$$\text{PRSM T LPA} / \text{PRSM USCT}$$

AV # OF SHRD LOG PROCS: The average number of logical processors operating in 'shared' mode on the CEC as derived in the following computation:

$$\text{PRSM T VP} / \text{PRSM USCT}$$

AV # OF ZAAP PROCESSORS: The average number of System z9 Application Assist Processors (zAAPs) actually assigned to the partition; derived from CA MICS data element LPCAVZAP (Avg Number of Active zAAP Processors). This data element is valid only for System z9 and later servers.

AV # OF ZIIP PROCESSORS: The average number of System z9 Integrated Information Processors (zIIPs) actually assigned to the partition; derived from CA MICS data element LPCAVSUP (Avg Number of Active zIIP Processors). This data element is valid only for System z9 and later servers.

AV # OF IFL PROCESSORS: The average number of integrated facilities for Linux actually assigned to the partition; derived from CA MICS data element LPCAVIFL (Avg Number of Active IFL Processors)

AV # ICF PROCESSORS: The average number of either internal coupling facilities (for System z9 and later) or special purpose processors (for previous systems) actually assigned to the partition; derived from CA MICS data element LPCAVICF (Avg Number of Active ICF Processors).

OF INTVLS PR/SM USED: The count of RMF recording intervals for which PR/SM was loaded: PRSMUSCT.

TOTAL RMF INTERVALS: The total number of RMF recording intervals in the reported time frame. This includes both intervals when the CEC was operating in PR/SM mode and intervals when it was not: INTERVLS.

Your_Company_Name
CA MICS Performance Manager
MONTHLY CEC SUMMARY REPORT
SYSTEM = ES01 LPARNAME = MVSESA01 RUNDATE = SEP 19, yyyy

MONTHS (YY-MM):	06-08	06-09	CHANGE %

UTILIZATION MEASURES			

CPU TM SHARD OR DED	1949246	1773736	-9.00 %
CPU TIME AVAILABLE	2235600	2041200	-8.69 %
PERCENT CEC CONSUMED	87.19	86.90	-0.33 %
ZAAP TM SHARD OR DED	116256	147243	1.27 %
ZAAP TIME AVAILABLE	2419202	2419202	-0.00 %
PERCENT ZAAP CONSUMED	4.81	4.10	-14.76 %
ZIIP TM SHARD OR DED	122.79	18.25	-85.14 %
ZIIP TIME AVAILABLE	108000	108000	---
PERCENT ZIIP CONSUMED	3.41	0.51	-85.04 %

ENVIRONMENT			

# OF PHYS PROCESSORS	3.00	3.00	---
AV # OF DED LOG PROCS	0.00	0.00	N/A
AV # DEFINED LOG PARTS	3.00	3.00	---
AV # OF ACTIVE PARTS	3.00	3.00	---
AV # OF SHRD LOG PROCS	5.00	5.00	---
AV # OF ZAAP PROCESSORS	1.00	1.00	---
AV # OF ZIIP PROCESSORS	1.00	1.00	---
AV # OF IFL PROCESSORS	1.00	1.00	---
AV # OF ICF PROCESSORS	2.00	2.00	---
# OF INTVLS PR/SM USED	828.00	756.00	-8.69 %
TOTAL RMF INTERVALS	828.00	756.00	-8.69 %

Figure 2-191. Sample Monthly CEC Summary Report

Your_Company_Name CA MICS Performance Manager HOURLY CEC SUMMARY REPORT SYSTEM = ES01 LPARNAME = MVSESA01 RUNDTE = Sep 16, yyyy										
HOURLY DATE (HR-MM/DD):	08-09/28	09-09/28	10-09/28	11-09/28	12-09/28	13-09/28	14-09/28	15-09/28	16-09/28	CHANGE %
UTILIZATION MEASURES										
CPU TM SHARD OR DED	6109.37	7226.77	9566.55	8101.38	5010.15	9435.97	10357.60	9983.69	7836.47	28.26 %
CPU TIME AVAILABLE	10799.99	10799.98	10799.98	10799.98	10799.98	10799.98	10799.98	10799.98	10799.98	-0.00 %
PERCENT CEC CONSUMED	56.57	66.91	88.58	75.01	46.39	87.37	95.90	92.44	72.56	28.26 %
ZAAP TM SHARD OR DED	17712	15984	16545	18489	28036	31752	49939	57456	64800	365.85 %
ZAAP TIME AVAILABLE	86400	86400	86400	86400	86400	86400	86400	86400	86400	---
PERCENT ZAAP CONSUMED	20.50	18.50	19.15	21.40	32.45	36.75	57.80	66.50	75.00	365.85 %
ZIIP TM SHARD OR DED	20.36	20.51	20.52	20.87	24.55	20.36	20.51	20.52	20.87	20.58 %
ZIIP TIME AVAILABLE	3600.00	3600.00	3600.00	3600.00	3600.00	3600.00	3600.00	3600.00	3600.00	---
PERCENT ZIIP CONSUMED	0.57	0.57	0.57	0.58	0.68	0.57	0.57	0.57	0.58	19.30 %
ENVIRONMENT										
# OF PHYS PROCESSORS	3.00	3.00	3.00	3.00	3.00	3.00	3.00	3.00	3.00	---
AV # OF DED LOG PROCS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
AV # DEFINED LOG PARTS	3.00	3.00	3.00	3.00	3.00	3.00	3.00	3.00	3.00	---
AV # OF ACTIVE PARTS	3.00	3.00	3.00	3.00	3.00	3.00	3.00	3.00	3.00	---
AV # OF SHRD LOG PROCS	5.00	5.00	5.00	5.00	5.00	5.00	5.00	5.00	5.00	---
AV # OF ZAAP PROCESSORS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	---
AV # OF ZIIP PROCESSORS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	---
AV # OF IFL PROCESSORS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	---
AV # OF ICF PROCESSORS	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	---
# OF INTVLS PR/SM USED	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	---
TOTAL RMF INTERVALS	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	---

Figure 2-192. Sample Hourly CEC Summary Report

2.6.3.3 LPAR Activity Report

The LPAR Activity Report tracks the utilization and environmental configuration information for one or more user selected logical partitions at a monthly, weekly, daily, or hourly level. Figure 2-193 shows an example of a monthly LPAR Activity Report and Figure 2-194 shows the hourly form of the report. In the hourly form, data for one day is displayed by hour using the DAYS timespan of the PR/SM LPAR Configuration and Activity File (HARLPC).

REPORT OPTIONS

The LPAR Activity Report supports standard time, SYSID, LPAR, and external data selection. These selections are controlled by the options discussed in Section 2.6.6.

DATA ELEMENTS

The elements reported are divided into the following categories:

- o Utilization Measures
- o Environment

UTILIZATION MEASURES

Utilization is shown in terms of the percentage of the total CEC consumed by the reported LPAR, and the percentages of shared and logical processor utilization. Items reported are:

CEC CPU The total available CPU time across the CEC,
TIME AVAIL: including both dedicated and shared processors
 as derived by the following computation:

$$(\text{PRSMTPP} / \text{INTERVALS}) * \text{DURATION}$$

LOG PROC The total of the processor dispatch time for all
DISPATCH TM: the logical processors assigned to the LPAR.
 When the Wait State Assist feature is disabled,
 this is the time that the LPAR was actually
 running on physical processors. When the Wait
 State Assist feature is enabled, this includes
 both the active time and the idle time the CP
 spends waiting for PR/SM time slices to expire:
 LPCTODTM.

PCT CEC CONSUMED BY LPR:	The percent of the total CEC consumed by this LPAR. It is computed as the ratio of the two preceding report items times 100.
PCT LPAR SHARED UTIL:	The percent of the shared processor resources used by this PR/SM logical partition (LPAR). This measurement does not apply to logical partitions with dedicated processors, and will have a value of zero: LPCPCSSU.
PCT LOGICAL PROC UTIL:	The percent of time that the LPAR's logical processors are in use for this LPAR. For LPARs with Wait Assist disabled, this is the percent CPU busy for the LPAR. For LPARs with Wait Assist enabled, this represents the percent of time that the processors are allocated to the LPAR. (When Wait Assist is enabled, LPCTODTM, the total processor dispatch time on which LPCPCVPU is based, includes time when the LPAR has put the processor in the wait state because there is no work for the processor to do.): LPCPCVPU.
PCT PROC EFFECTIVE DISP:	The percentage of total dispatch time that was effectively used by LPARs, exclusive of LPAR management time. It is derived by the following computation: $\text{PRSMTEDT} / \text{PRSMTTDT} * 100$
PCT PR/SM OVERHEAD:	The percentage of total dispatch time that was attributable to LPAR management functions for all LPARs in the CEC. This includes time used in support of individual partitions, as well as unattributable processor time. It is derived by the following computation: $\text{PRSMTLMT} / \text{PRSMTTDT} * 100$
LPAR ZAAP PROC UTILIZATION:	The percentage of time that the System z9 Application Assist Processors (zAAPs) are in use for this logical partition; derived from CA MICS data element LPCPCVZU (LPAR zAAP Processors Utilization). This data element is valid only for System z9 and later servers.
LPAR TOTAL ZAAP PROC UTILIZATION:	The percentage of time that the System z9 Application Assist Processors (zAAPs) were dispatched to the PR/SM logical partition;

derived from CA MICS data element LPCPCTZU (LPAR Total zAAP Processors Utilization). For shared LPARs with Wait Assist enabled or for dedicated LPARs, LPCPCTZU includes both busy time and wait time. For shared LPARs with Wait Assist disabled, this element is equal to the percentage of shared zAAP resources used by this LPAR. This element is valid only for System z9 and later servers.

LPAR ZIIP PROC UTILIZATION: The percentage of time that the System z9 Integrated Information Processors (zIIPs) are in use for this logical partition; derived from CA MICS data element LPCPCVSU (LPAR zIIP Processors Utilization). This data element is valid only for System z9 and later servers.

LPAR TOTAL ZIIP PROC UTILIZATION: The percentage of time that the System z9 Integrated Information Processors (zIIPs) were dispatched to the PR/SM logical partition; derived from CA MICS data element LPCPCTDU (LPAR Total zIIP Processors Utilization). For shared LPARs with Wait Assist enabled or for dedicated LPARs, LPCPCTDU includes both busy time and wait time. For shared LPARs with Wait Assist disabled, this element is equal to the percentage of shared zIIP resources used by this LPAR. This element is valid only for System z9 and later servers.

ENVIRONMENT

The environmental measures appearing in this report reflect the LPAR configuration and monitoring activity.

AV # OF LOG PROCESSORS: The average number of logical processors assigned to this LPAR as derived by the following computation:

$$\text{LPCNOCPU} / \text{LPCTOACT}$$

AV # OF ZAAP PROCESSORS: The average number of System z9 Application Assist Processors (zAAPs) actually assigned to the partition; derived from CA MICS data element LPCAVZAP (Avg Number of Active zAAP Processors). This data element is valid only for System z9 and later servers.

AV # OF ZIIP The average number of System z9 Integrated

PROCESSORS: Information Processors (zIIPs) actually assigned to the partition; derived from CA MICS data element LPCAVSUP (Avg Number of Active zIIP Processors). This data element is valid only for System z9 and later servers.

AV # OF IFL PROCESSORS: The average number of integrated facilities for Linux actually assigned to the partition; derived from CA MICS data element LPCAVIFL (Avg Number of Active IFL Processors)

AV # ICF PROCESSORS: The average number of either internal coupling facilities (for System z9 and later) or special purpose processors (for previous systems) actually assigned to the partition; derived from CA MICS data element LPCAVICF (Avg Number of Active ICF Processors).

OF INTVLS LPAR ACTIVE: The count of RMF recording intervals when the reported LPAR was active: LPCTOACT.

INTVLS W/WSA ENABLED: The number of RMF recording intervals in which the LPAR Wait State Assist feature was enabled: LPCWSECT.

INTVLS W/WSA DISABLED: The number of RMF recording intervals in which the LPAR Wait State Assist feature was disabled: LPCWSDCT.

INTVLS W/DED PROCESSOR: The number of RMF recording intervals when the reported LPAR had one or more dedicated processors while the CEC was operating in PR/SM mode. These are processors defined to be used only by this LPAR: LPCDEDCT.

TOTAL RMF INTERVALS: The total number of RMF recording intervals in the reported time frame. This includes both intervals when the CEC was operating in PR/SM mode, and intervals when it was not: INTERVLS.

```

                                CA MICS Performance Manager
                                MONTHLY LPAR ACTIVITY REPORT
                                FOR ZONE = 1
SYSTEM = ES01  LPARNAME = MVSESA01  RUNDATE =  SEP 19, yyyy
-----
| MONTHS (YY-MM):          yy-08    yy-09    CHANGE %  |
|-----|

```

UTILIZATION MEASURES			

CEC CPU TIME AVAIL	2235600	2041200	-8.69 %
LOG PROC DISPATCH TM	1065710	1098778	3.10 %
PCT CEC CONSUMED BY LPR	47.67	53.83	12.92 %
PCT PR/SM OVERHEAD	36.32	19.51	-46.28 %
PCT PROC EFFECTIVE DISP	63.68	80.49	26.39 %
PCT LPAR SHARED UTIL	47.67	53.83	12.92 %
PCT LOGICAL PROC UTIL	47.67	53.83	12.92 %
LPAR ZAAP PROC UTILIZATION	11.00	11.00	---
LPAR TOTAL ZAAP PROC UTIL	11.00	11.00	---
LPAR ZIIP PROC UTILIZATION	14.00	14.00	---
LPAR TOTAL ZIIP PROC UTIL	14.00	14.00	---

ENVIRONMENT			

AV # OF LOG PROCESSORS	3.00	3.00	---
AV # OF ZAAP PROCESSORS	1.00	1.00	---
AV # OF ZIIP PROCESSORS	1.00	1.00	---
AV # OF IFL PROCESSORS	1.00	1.00	---
AV # OF ICF PROCESSORS	2.00	2.00	---
# OF INTVLS LPAR ACTIVE	828.00	756.00	-8.69 %
INTVLS W/WSA ENABLED	0.00	0.00	N/A
INTVLS W/WSA DISABLED	828.00	756.00	-8.69 %
INTVLS W/DED PROCESSOR	0.00	0.00	N/A
TOTAL RMF INTERVALS	828.00	756.00	-8.69 %

Figure 2-193. Sample Monthly LPAR Activity Report

CA MICS Performance Manager HOURLY LPAR ACTIVITY REPORT SYSTEM = ES01 LPARNAME = MVESA01 RUNDATE = SEP 19, yyyy										
HOUR-DATE (HR-MM/DD):	08-09/28	09-09/28	10-09/28	11-09/28	12-09/28	13-09/28	14-09/28	15-09/28	16-09/28	CHANGE %

UTILIZATION MEASURES										

CEC CPU TIME AVAIL	10799.99	10799.98	10799.98	10799.98	10799.98	10799.98	10799.98	10799.98	10799.98	-0.00 %
LOG PROC DISPATCH TM	5280.15	6348.71	6591.22	6172.18	4049.98	5420.50	7815.93	7011.34	5102.98	-3.35 %
PCT CEC CONSUMED BY LPR	48.89	58.78	61.03	57.15	37.50	50.19	72.37	64.92	47.25	-3.35 %
PCT PR/SM OVERHEAD	36.32	24.14	28.00	25.35	21.89	23.91	18.29	20.59	19.51	-46.28 %
PCT PROC EFFECTIVE DISP	63.68	75.86	72.00	74.65	78.11	76.09	81.71	79.41	80.49	26.39 %
PCT LPAR SHARED UTIL	48.89	58.78	61.03	57.15	37.50	50.19	72.37	64.92	47.25	-3.35 %
PCT LOGICAL PROC UTIL	48.89	58.78	61.03	57.15	37.50	50.19	72.37	64.92	47.25	-3.35 %
LPAR ZAAP PROC UTILIZATION	0.00	0.00	0.00	0.00	0.00	0.00	8.45	8.45	8.45	N/A
LPAR TOTAL ZAAP PROC UTIL	0.00	0.00	0.00	0.00	0.00	0.00	8.45	8.45	8.45	N/A
LPAR ZIIP PROC UTILIZATION	0.21	0.22	0.22	0.22	0.21	0.22	0.22	0.22	0.23	9.53 %
LPAR TOTAL ZIIP PROC UTIL	0.21	0.22	0.22	0.22	0.21	0.22	0.22	0.22	0.23	9.53 %

ENVIRONMENT										

AV # OF LOG PROCESSORS	3.00	3.00	3.00	3.00	3.00	3.00	3.00	3.00	3.00	---
AV # OF ZAAP PROCESSORS	0.00	0.00	0.00	0.00	0.00	0.00	1.00	1.00	1.00	N/A
AV # OF ZIIP PROCESSORS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	N/A
AV # OF IFL PROCESSORS	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	---
AV # OF ICF PROCESSORS	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	---
# OF INTVLS LPAR ACTIVE	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	---
INTVLS W/WSA ENABLED	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
INTVLS W/WSA DISABLED	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	---
INTVLS W/DED PROCESSOR	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
TOTAL RMF INTERVALS	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	4.00	---

Figure 2-194. Sample Hourly LPAR Activity Report

2.6.4 Usage Guidelines

This section helps you use the PR/SM LPAR performance analysis functions of the CA MICS Performance Manager Option. The following topics are discussed:

- 1 - CEC Configuration Tuning
- 2 - LPAR Interaction Analysis
- 3 - Using the Wait State Assist Feature

2.6.4.1 CEC Configuration Tuning

In tuning the CEC configuration, you should be aware of high levels of contention for shared processor resources. The overhead required for PR/SM to manage and allocate processor resources on behalf of the logical partitions can grow to as much as 12% as more processors are shared and as contention for their use increases.

The PR/SM CEC Summary report described in section 2.6.3.2 provides a key indicator of PR/SM overhead in the form of the Percent of CEC Consumed. This measure is particularly useful when all physical processors are being shared, as it then reflects the overall utilization of processor resources in the complex. However, if dedicated processors are included in the configuration, each such processor is considered as 100% utilized (since it is not available to the CEC for sharing with other logical partitions (LPARs)).

Significant benefits can be derived from the use of dedicated processors. If multiple workloads that previously ran simultaneously on a multi-processor system can be re-distributed such that they run separately on multiple LPARs each with a dedicated CP, increased throughput is possible due to the reduced overhead that occurs because a single SCP is not trying to juggle multiple processors.

2.6.4.2 LPAR Interaction Analysis

When executing in LPAR mode, MVS/XA or ESA can provide measurements that cover activity within the LPAR along with overall activity of the other active LPARs. With PR/SM, many, if not all, of the effects of sharing a physical complex among multiple SCPs can be characterized with measurement data from a single source.

Since memory and channel paths are dedicated to individual LPARs, analyzing resource utilization is largely unchanged when compared to non-PR/SM operations. The single exception is in the measurements of the activity of the I/O Processor (IOP) of a 3090 CEC. These are global to the IOP and, therefore, apply to the sum of the activity of all active LPARs. There is no direct way to determine the load on the IOP imposed by a particular LPAR.

As long as the demand for logical processor resources is less than the availability of CPs, the activity in one LPAR will have little or no impact on other LPARs. It is only during periods of contention for processor resources that inter-partition effects can be expected.

The most obvious effect of contention is that less work being done by the CPs than would have been done in a non-PR/SM complex because PR/SM prioritizes LPAR access to CPs based on the relative weights specified by the user for each active LPAR in the complex. The relative share of processor resources available to each LPAR can be controlled by modifying the relative weights.

The processing weights for each LPAR are set when the LPAR is defined, and may be modified for an active partition at any time by the console operator. Their actual numeric values are meaningful only when compared to each other. Before being applied, the weight is divided by the number of active logical processors for the LPAR. The weights are then applied to the sum of the activity of all logical processors active in the LPAR. Thus, during a period of high contention for processor resources, the sum of the dispatch times for all of the logical processors for an LPAR with a relative weight of 200 will be twice the sum of the dispatch times of a second LPAR with a weight of 100.

An LPAR with fewer logical processors defined will receive a higher proportion of an individual logical processor than an LPAR with more logical processors defined. An important consideration here is that the number of logical processors

defined for an LPAR can limit its total processor utilization to a value less than that implied by its relative weight. In other words, if an LPAR has only one logical processor defined, it can never receive more than 25 per cent of the resources of a 3090-400 even if it has a high relative weight as compared to the weights of any other active LPARs.

2.6.4.3 Using the Wait State Assist Feature

A key component of PR/SM is the Wait State Assist (WSA) feature. WSA allows an installation to define what happens in a shared processor environment when an LPAR goes into a wait state prior to the end of its time slice. When the feature is disabled (the usual setting), a CP will remain allocated to an LPAR until the time slice expires or the SCP running in the LPAR attempts to place the logical processor into a wait state; at this time, the CP will be made available for work elsewhere.

When WSA is enabled, PR/SM will not take control of a CP when the logical processor goes into a wait state, but leaves the CP under the control of the LPAR for the entire duration of the time slice. Ordinarily, fixed time slicing would be utilized with WSA enabled.

Since one of the primary objectives of PR/SM is to greatly reduce (if not eliminate) the time a logical processor is in a wait state, enabling WSA should be weighed carefully. Generally speaking, enabling WSA effectively disables one of PR/SM's key tools for attempting to enhance processor utilization. Nevertheless, there may be certain environments where it is desirable to allow individual SCPs to retain control of a CP, thereby preventing the overhead of PR/SM processor switching. This might be the case in highly CP intensive workloads. However, the use of dedicated processors might be more appropriate than utilizing WSA.

In any case, both the CEC Summary and LPAR Activity reports discussed in section 2.6.3 should be monitored for the effects of enabling or disabling WSA.

2.6.5 Reserved

2.6.6 PR/SM LPAR Analysis Operation

This section describes the operation of the PR/SM LPAR Analysis routines in terms of the CA MICS Performance Manager inquiry options selected.

Inquiry options available with the PR/SM LPAR Analysis software describe the reports to be generated. In particular, these options define:

- o The cycles of the CA MICS HARLPC file to be processed.
- o Workload selection by system, logical partition, date/time, and zone.
- o The summarization level of the reports as hourly, daily, weekly, or monthly. The availability of each summarization level assumes that your system supports the required cycles at each level.
- o The reports to be produced.

The following fields are common to most control panels:

- Timespan: The timespan (DETAIL, DAYS, WEEKS, or MONTHS) of the HARLPC file that you wish to use as input for processing.
- Cycle(s): One or more cycles of the CA MICS file that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a CYCLE from the CA MICS file will usually contain a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided by the DATE and HOUR control options.
- Database: The single-character CA MICS identifier for the database to be accessed. Valid values are the database unit identifiers available at your installation. Contact your CA MICS System Administrator if you do not know what database units are available, or enter a question mark (?) to view a list of your database units. You can select up to five databases per inquiry.

Allow Run-Time Override:	The indication of whether the database identifier and/or the cycles specification can be overridden at execution time.
Report by hour:	The indication of whether a report or plot should be generated for each hour.
Generate a report for each value of Zone:	The indication of whether a separate report or plot should be generated for each zone defined in the input data or whether all zones should be merged together for the report or plot.
Date:	Up to two date ranges (DDMONYY format) to be used for analysis. This option is normally used to drop the extra day from a cycle that is due to the timing of the SMF dump.
Hour:	One to three ranges of hours for which data is to be reported. For an individual hour, leave the end hour blank or make it the same as the start hour. You can specify a value from 00 to 23.
Zone:	The specification that zone selection is to be applied. If zone 1 definition is 8:00 a.m. to 6:00 p.m., you can specify reports covering that time period by entering a 1 in the ZONE field.
SYSID:	The specification that SYSIDs are to be selected for reporting. You can enter up to six SYSIDs to be selected. If you leave the field blank, no testing of SYSIDs will be done.
LPAR Name:	The logical partitions to be selected for reporting. If you code this field, only those specified will be selected.
Specify extended options:	The specification that allows you to supply additional report specifications. These options include special SAS data selection statements, non-standard input and output files, and special graphics options.

The operation of the time selection control parameters depends on definitions made during the installation of CA MICS. If you are not sure about the ZONE definitions for your installation or the CA MICS database update procedures, consult your CA MICS System Administrator.

Figure 2-195 shows the PR/SM LPAR (PRA) Analysis menu used to select the desired report within the PR/SM reporting area.

[illegible]

Figure 2-195. PR/SM LPAR Analysis Selection Panel

The following sections describe the inquiry panels for each PR/SM report:

- 1 - PR/SM CEC Summary Report
- 2 - PR/SM LPAR Activity Report

2.6.6.1 PR/SM CEC Summary Report

The PR/SM CEC Summary Report details the utilization and configuration of an entire Central Electronic Complex. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS time-spans if they are present in your CA MICS database. The PR/SM CEC Summary Report obtains its data from the HARLPC Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The PR/SM CEC Summary Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options.

```

----- PR/SM CEC Summary Report -----
|
| Command ==>
|
| Composing CA MICS Inquiry: PRSM01 - Track CEC Activity
|
| Time-span ==> MONTHS (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 02
| Data base ==> P (PRIMARY)      _      _      _      _
| Allow Run-Time Override of Data Base ==> Y (Y/N) and/or Cycle(s) ==> Y (Y/N)
|
| Report by hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone      ==> N (Y/N)
|
| Data Selection:                                           Allow run-time override
| Date      ==> _____ - _____ - _____ ==> N (Y/N)
| Hour      ==> __ - __      __ - __      __ - __      ==> N (Y/N)
| Zone      ==> _ - _ - _ - _ - _ - _ - _ - _ - _ - _ ==> N (Y/N)
| Sysid     ==> _____ ==> N (Y/N)
| LPAR Name  ==> _____
|
|
| Specify extended options ==> N (Y/N)
|
-----/

```

Figure 2-196. PR/SM CEC Summary Report Inquiry Panel

Figure 2-196 shows a PR/SM CEC Summary Report inquiry panel. The fields common to all PR/SM Analysis routines are described in Section 2.6.6. The following are general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Report by hour and Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized to the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

Duplicate information about an entire CEC appears in the HARLPC file for each LPAR that is defined on the complex. Additionally, if more than one SYSID is generating RMF data for the same CEC, duplicate information is retained for each SYSID/LPAR combination. Consequently, to prevent having duplicate reports, you should determine one SYSID/LPAR combination for each CEC that you wish to have reported.

If multiple SYSIDs have created RMF data for the same complex (for example, two different MVS systems are running in separate logical partitions on the same CEC), then the most comprehensive reporting will be achieved by selecting the SYSID that was most active during the time period to be reported.

Four other panels in this series are described in this section that you can use to specify additional options.

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.6.6.2 PR/SM LPAR Activity Report

The PR/SM LPAR Activity Report displays the utilization and configuration of a logical partition operating in a Central Electronic Complex. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the PR/SM LPAR Activity Report obtains its data from the HARLPC Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The PR/SM LPAR Activity Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options.

```

/----- PR/SM LPAR Activity Report ----- \
| Command ==>                                     |
| Composing CA MICS Inquiry: PRSM02 - Track LPAR Activity |
| Time-span ==> MONTHS (DETAIL/DAYS/WEEKS/MONTHS)      Cycle(s) ==> 01 - 02 |
| Data base ==> P (PRIMARY) _                         _ |
| Allow Run-Time Override of Data Base ==> Y (Y/N) and/or Cycle(s) ==> Y (Y/N) |
| Report by hour if timespan is DAYS or WEEKS ==> N (Y/N) |
| Generate a report for each value of Zone ==> N (Y/N) |
| Data Selection:                                     Allow run-time override |
| Date ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
| Hour ==> _ _ - _ _ - _ _ - _ _ - _ _ ==> N (Y/N) |
| Zone ==> _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
| Sysid ==> _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N) |
| LPAR Name ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-197. PR/SM LPAR Activity Report Inquiry Panel

Figure 2-197 shows a PR/SM CEC Summary Report inquiry panel. The fields common to all PR/SM Analysis routines are described in Section 2.6.6. The following are general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS time spans. Summarization can be specified with the Report by hour and Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization

parameter, the report is summarized to the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

If no LPAR name is selected, a report will be generated for each LPAR on each selected SYSID (or all SYSIDs if none are specified). If more than one SYSID is generating RMF data for the same CEC, duplicate information is retained for each SYSID/LPAR combination. Consequently, to prevent having duplicate reports created, you should determine one SYSID/LPAR combination for the logical partition you wish to have reported.

If multiple SYSIDs have created RMF data for the same complex (for example, two different MVS systems are running in separate logical partitions on the same CEC), then the most comprehensive reporting will be achieved by selecting the SYSID that was most active during the time period to be reported.

There are four other panels in this series that you can use to specify additional options. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.6.6.2 PR/SM LPAR Activity Report

The PR/SM LPAR Activity Report displays the utilization and configuration of a logical partition operating in a Central Electronic Complex. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the PR/SM LPAR Activity Report obtains its data from the HARLPC Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The PR/SM LPAR Activity Report is produced using a series of five panels. Only the first panel is required. The others are used to specify additional options.

```

/----- PR/SM LPAR Activity Report -----\
| Command ==>
|
| Composing CA MICS Inquiry: PRSM02 - Track LPAR Activity
|
| Time-span ==> MONTHS (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 02
| Data base ==> P (PRIMARY)      _      _      _      _
| Allow Run-Time Override of Data Base ==> Y (Y/N) and/or Cycle(s) ==> Y (Y/N)
|
| Report by hour if timespan is DAYS or WEEKS ==> N (Y/N)
| Generate a report for each value of Zone      ==> N (Y/N)
|
| Data Selection:                                           Allow run-time override
|   Date      ==> _____ - _____ - _____ ==> N (Y/N)
|   Hour      ==> __ - __      __ - __      __ - __      ==> N (Y/N)
|   Zone      ==> _ _ _ _ _ _ _ _ _ _ ==> N (Y/N)
|   Sysid     ==> _____ ==> N (Y/N)
|   LPAR Name ==> _____
|
|
| Specify extended options ==> N (Y/N)
\-----/

```

Figure 2-197. PR/SM LPAR Activity Report Inquiry Panel

Figure 2-197 shows a PR/SM CEC Summary Report inquiry panel. The fields common to all PR/SM Analysis routines are described in Section 2.6.6. The following are general comments on this panel.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS time spans. Summarization can be specified with the Report by hour and Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized to the timespan level (i.e., MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

If no LPAR name is selected, a report will be generated for each LPAR on each selected SYSID (or all SYSIDs if none are specified). If more than one SYSID is generating RMF data for the same CEC, duplicate information is retained for each SYSID/LPAR combination. Consequently, to prevent having duplicate reports created, you should determine one SYSID/LPAR combination for the logical partition you wish to have reported.

If multiple SYSIDs have created RMF data for the same complex (for example, two different MVS systems are running in separate logical partitions on the same CEC), then the most comprehensive reporting will be achieved by selecting the SYSID that was most active during the time period to be reported.

There are four other panels in this series that you can use to specify additional options. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.6.7 Requirements

CA MICS PR/SM reports enable you to track the utilization, configuration and monitoring activity of complexes and logical partitions on a monthly, weekly, daily, or hourly level. Utilization measures are reported in terms of the percentage of a CEC consumed and shared and logical processor usage. Environmental measures are also provided to allow for the tracking of configuration changes over time, and their impact on system usage.

Figure 2-198 shows the CA MICS files and timespan levels that must be maintained at your installation in order to produce all of the PR/SM LPAR Analysis reports.

CA MICS	LEVEL
FILE	-----
	X D W M Y

HARLPC	x x x x

Figure 2-198. PR/SM LPAR Analysis CA MICS File Requirements

For these reports, only one level of the HARLPC file is required (i.e., DETAIL, DAYS, WEEKS, or MONTHS), depending on the timespan you specify during inquiry definition.

The number of cycles that are maintained at your installation for each of the CA MICS files in Figure 2-198 determines the time frame for which you may produce PR/SM LPAR Analysis reports.

2.7 Graphic Analysis

This section describes the graphic analyses available in the Performance Manager. Graphic analysis of performance and service indicators provides you with several series of graphs to assist you in tracking and analyzing system behavior and key service related indicators. This chapter provides guidelines for interpreting these graphs and for producing them. These graphs can provide you views of system and workload performance for a variety of time perspectives from the very detailed to a perspective covering weeks or months.

The CPU graphs can help you understand both the utilization patterns of your processors and the distribution of this utilization among various workloads. This can help you ensure that the processors, as critical resources are being used to satisfy your key workloads.

The storage graphs show you the distribution of storage among the various system areas as well as the distribution of storage use among workloads. As with the CPU graphs, these storage graphs help you ensure that storage, as a critical resource, is being used to satisfy your key workloads.

The TSO graphs can assist you in tracking MVS response time for TSO and easily focuses your attention toward periods during which service levels were exceeding expectations or agreed upon objectives.

The I/O tracking graphs can help you focus on particular disk devices and time periods that may have experienced I/O bottlenecks.

The paging graphs help you to manage the performance of the MVS auxiliary storage. They will help you quickly discover whether any potential performance problems exist with your paging volumes and assist you in focusing rapidly on any problem when it is discovered.

Producing graphs is discussed in the final sections of this chapter. These sections can assist you in choosing graphic options, the time periods and systems upon which to report, and the methods for providing workload descriptions.

The following sections in this chapter discuss the use and interpretation of these graphs:

- 1 - CPU Graphs
- 2 - Storage Graphs

- 3 - TS0 Service Graphs
- 4 - I/O Tracking Graphs
- 5 - Paging Graphs
- 6 - Producing Graphs

2.7.1 CPU Graphs

In managing performance, it is important to understand and track two very basic types of information about your processors: how busy they are and what workloads are using these resources. z/OS systems should be able to operate effectively at very high processor utilization, assuming no other system resource is limiting critical response time. You need to know which z/OS service classes and workloads account for the majority of this utilization.

These utilization values cannot be viewed in isolation. If utilization is high but no workloads spend time waiting for a processor, and workloads are experiencing acceptable online response time or batch turnaround time, there may be no need for further diagnosis and corrective action. The use of other indicators along with CPU busy can help you determine whether there is a processor utilization problem or some other form of workload delay.

This section discusses the following graphic analyses:

- 1 - CPU Busy by Hour and Day
- 2 - CPU Busy by Service Class
- 3 - CPU Busy and Other Indicators
- 4 - CPU Busy by Workload Type - Pie Chart
- 5 - CPU Busy by System
- 6 - CPU Busy by Workload Type - Bar Chart
- 7 - Workload Velocity and CPU Activity

2.7.1.1 CPU Busy by Hour and Day

GRAPHIC RESULTS

This graphic analysis provides a multiline plot of CPU busy, plotted by hour of the day, with each line representing data from a different day.

ANALYSIS GUIDELINES

This plot shows you CPU busy information for a specific system for multiple days. You should look here to ensure that the data being analyzed is both repeatable and consistent. You should know whether you are analyzing a typical day, a peak load day that is to be expected with some regularity, or an unusual occurrence that might be repeated only very rarely.

Notice the CPU busy patterns throughout a day. When analyzing performance, it is important to know the characteristic times of peak loads during the day. These patterns also allow you to readily see when there are CPU resources available. If you know the major workloads that run on a processor, you may be able to determine when these workloads are most significant to overall utilization. This is determined by noting the times when these workloads normally begin to increase and taper off. Normally, when looking for CPU problems, you would select the times when the systems are running at maximum utilization.

The detailed behavior of workloads and the system influences that may be affecting them can be further analyzed using the Installation Performance Specification analysis reports. Note that, when in goal mode, the IPS controls are "hidden" so you can't change them. The Workload Manager decides the settings and automatically implements them.

2.7.1.2 CPU Busy by Service Class

GRAPHIC RESULTS

This graphic analysis provides a stacked vertical bar chart of CPU utilization by one or more Workload Manager service classes. You can display all active classes on the plot or you can select several classes that are of particular interest to you. Overhead not measured as part of any service class is spread evenly across all groups or classes for reporting.

ANALYSIS GUIDELINES

Similar considerations come into play for environments that are operating in "goal mode." In this case, workload granularity must be established through the definition of service classes using the Workload Manager application. The Workload Manager is used to define both service classes that dictate distribution of resources (Standard Service Classes) and service classes designed for simply reporting utilization and service measures (Report Service Classes).

Note: Having multiple applications running in the same service class will restrict your ability to distinctly analyze and manage the work on your system.

The detailed behavior of workloads and the system influences that may be affecting them can be further analyzed using the Installation Performance Specification analysis reports.

2.7.1.3 CPU Busy and Other Indicators

GRAPHIC RESULTS

This graphic analysis provides a multiline plot of CPU busy and several other indicators that will be non-zero whenever work was ready to use the processor but was not dispatched. In addition, paging activity is displayed, because it is another potential reason for slow processing.

ANALYSIS GUIDELINES

In addition to CPU busy, there are other indicators found on the graph that will help you understand if there are workload delays.

OCPU represents the percentage of time that all engines were dispatched and there was still at least one address space waiting to be dispatched. O2 and O3 are the same except they represent at least 2 and 3 or more address spaces waiting to be dispatched.

OREADY is the number of address spaces swapped out and waiting to be swapped in. They are ready to do work, but, due to the SRM constraints, they are kept swapped out.

The final indicator, PAGEIN, shows the rate at which pages are being moved from auxiliary storage to central storage. As an indicator of contention for real storage, it can show you whether inadequate storage may be a contributing cause of sluggish performance.

This chart can show whether latent demand exists: whether work is being delayed that was otherwise ready to proceed. It does not show which applications are experiencing or causing contention.

The detailed behavior of workloads and the system influences that may be affecting them can be further analyzed using the Installation Performance Specification analysis reports.

2.7.1.4 CPU Busy by Workload Type - Pie Chart

GRAPHIC RESULTS

This graphic analysis provides a pie chart of the distribution of processor utilization among major workloads. Ordinarily you would describe the rules for grouping particular service classes into certain workloads. An example of the method for defining workloads is available in Section 2.7.6.3. If you do not provide such a method, each service class will be displayed as a separate workload.

ANALYSIS GUIDELINES

This chart presents a high level overview for a major system resource, CPU capacity. With it you can show where your leverage to manage utilization is by workload type. Shifting a workload to another processor might not relieve a processor bottleneck if its relative percentage of the total utilization is small. Because this graph is a pie chart, the displayed CPU utilization represents the percentage of total CPU busy caused by each workload type. This displayed utilization is the relative share of CPU and not the total amount.

The detailed behavior of workloads and the system influences that may be affecting them can be further analyzed using the sysplex analysis reports described in Section 2.8.

SPECIAL CONSIDERATIONS

If you do not provide a workload description, as shown in Section 2.7.6.3, each service class will be treated as a separate workload. This may provide so many unique categories that SAS/GRAPH will shift from using a pie chart to using a horizontal bar chart. While a horizontal bar chart prevents the loss of information from the chart, you will probably find analysis and reporting more meaningful if you describe workloads.

2.7.1.5 CPU Busy by System

GRAPHIC RESULTS

This graphic analysis provides a multiline plot of CPU busy, plotted over time, with each line representing data from a different system (SYSID).

ANALYSIS GUIDELINES

This plot shows you CPU busy information for a specific time period (day, week, month, etc.) for multiple systems. You should look here to ensure that the data being analyzed is both repeatable and consistent. From this graph, you can see whether you are analyzing a typical time period, a peak load period (for example, prime shift weekdays) that is to be expected with some regularity, or an unusual occurrence that might be repeated only very rarely.

Notice the CPU busy patterns throughout a day, especially contrasting one system with another. When analyzing performance, it is important to know the characteristic times of peak loads during the day by system. These patterns also allow you to readily see when there are CPU resources available. If you know the major workloads that run on each processor, you may be able to determine when these workloads are most significant to overall utilization, and also how workloads might be shifted from one system to another.

The detailed behavior of workloads and the system influences that may be affecting them can be further analyzed.

2.7.1.6 CPU Busy by Workload Type - Bar Chart

GRAPHIC RESULTS

This graphic analysis provides a vertical bar chart of the distribution of processor utilization among major workloads. Ordinarily you would describe the rules for grouping particular service classes into certain workloads. An example of the method for defining workloads is available in Section 2.7.6.3. If you do not provide such a method, each service class will be displayed as a separate workload.

ANALYSIS GUIDELINES

As with the corresponding pie chart described in Section 2.7.1.4, this chart presents a high-level overview for a major system resource, CPU capacity. With it you can show where your leverage to manage utilization is by workload type. For example, shifting a workload to another processor might not relieve a processor bottleneck if its relative percentage of the total utilization is small.

Because this graph is a bar chart, the displayed CPU utilization represents the total utilization segmented by each workload type. This displayed utilization includes the uncaptured CPU utilization allocated proportionately to each workload.

The detailed behavior of workloads and the system influences that may be affecting them can be further analyzed using the sysplex analysis reports described in Section 2.8.

SPECIAL CONSIDERATIONS

If you do not provide a workload description, as shown in Section 2.7.6.3, each service class will be treated as a separate workload. Further, if you choose to use "report service classes" (rather than standard classes), overlapping resource utilization will tend to overstate actual utilization. Therefore, this should only be done with a full understanding of the impact on the resulting analysis.

2.7.1.7 Workload Velocity and CPU Activity

GRAPHIC RESULTS

The Workload Velocity and CPU Activity Graph provides a snapshot of the performance of workloads that have been defined with service goal type (WLMGOLTP) of "Execution of Velocity". The Goal (WLMGOLVE) and actual velocity (SECEXVEL) is expressed as a percentage. Actual velocity is a measure of how fast the transactions with this goal type are running. Work that is "not ready" is not included in this percentage. The CPU utilization, expressed as a percentage, is available for interval comparisons with the given workload.

ANALYSIS GUIDELINES

The reports provide a method to verify WLM execution goal success. The user may note that systems which are running CPU constrained during certain hours of the day may be unable to achieve their WLM Goal. This report will identify the constrained/peak hours.

SPECIAL CONSIDERATIONS

CPU is just one component of performance and some factors impact each other (i.e., I/O affects CPU utilization). Some factors to consider are: I/O wait (including data set contention), coupling facility wait, tape mount wait (rewind time, SILO robot response, computer operator mount wait) and I/O error recovery. To understand fully the performance of workloads, all contributing factors must be accounted for and understood.

Figure 2-199 provides an example of the Workload Velocity and CPU line graph.




Figure 2-199. Sample Workload Velocity and CPU Activity Graph

2.7.2 Storage Graphs

Managing storage can help ensure that online response and batch throughput service remain acceptable. Overcommitment of total storage frames can lead to dramatic degradation in system performance as MVS responds to contention for central and expanded storage frames by increasing demand paging involving auxiliary storage on disk devices.

These storage graphs will help you understand and report on the usage of combined central and expanded storage frames by the various system areas of MVS and the private area. The second graph will help you understand the distribution of active usage of storage among various workloads.

This section discusses the following graphic analyses:

- 1 - System Storage Map
- 2 - Storage Used by Workload Type-Bar Chart
- 3 - Storage Used by Performance Group/Service Class
- 4 - Storage Used by Workload Type-Pie Chart

2.7.2.1 System Storage Map

GRAPHIC RESULTS

This graphic analysis provides a vertical bar chart of storage frame use by the major system areas and the private area.

ANALYSIS GUIDELINES

This graph maps total storage frames allocated by MVS system areas. This chart shows the storage frames allocated to the Nucleus, CSA, LPA, SQA, LSQA, Private, and Available frames. The sum of these fields represents the total number of central storage frames. When this resource becomes constrained, MVS will respond with increased paging and increased swapping to auxiliary storage.

Frame allocation is a type of resource usage that can be analyzed effectively from a variety of perspectives. Long range charts done using the CA MICS MONTHS timespan can permit you to track trends in frame usage. More detailed charts done using the DAYS timespan can help you analyze peak period loads to determine if frame contention might be a factor in any performance degradation.

You can analyze memory use further using the Real Storage Management analysis reports described in Section 2.2. Virtual storage allocation patterns can have an effect upon processor storage usage. These can be analyzed using the Virtual Storage Management analysis reports found in Section 2.5.

2.7.2.2 Storage Used by Workload Types-Bar Chart

GRAPHIC RESULTS

This graphic analysis provides a vertical bar chart of total storage frames used by various defined workloads.

ANALYSIS GUIDELINES

Total storage includes both central and expanded frames. This chart can help you understand what workloads are your major consumers of storage frames. It can also help you determine if storage usage by a competing workload appears correlated with a period when other measurements reported during that same period were degraded.

You can analyze memory use further using the Real Storage Management analysis reports described in Section 2.2. Virtual storage allocation patterns can have an effect upon processor storage usage. These can be analyzed using the Virtual Storage Management analysis reports found in Section 2.5. Each of these chapters can help you understand and manage memory use by particular workloads.

SPECIAL CONSIDERATIONS

If you do not provide a workload description, each performance group or service class (for a system operating in "goal mode") will be treated as a separate workload. This may provide so many unique categories that your analysis may prove difficult. You will probably find analysis and reporting more meaningful if you describe workloads.

2.7.2.3 Storage Used by Performance Group/Svc Class

GRAPHIC RESULTS

This graphic analysis provides a vertical bar chart of total storage frames used by all or selected performance groups or service classes.

ANALYSIS GUIDELINES

Total storage includes both central and expanded frames. This chart can help you understand what performance groups or service classes are your major consumers of storage frames. It can also help you determine if storage usage by a competing group or class appears correlated with a period when other measurements reported during that same period were degraded.

You can analyze memory use further using the Real Storage Management analysis reports described in Section 2.2. Virtual storage allocation patterns can have an effect on processor storage usage. These can be analyzed using the Virtual Storage Management analysis reports found in Section 2.5. Each of these chapters can help you understand and manage memory use by particular workloads.

SPECIAL CONSIDERATIONS

If there are a lot of performance groups or service classes defined, there may be so many unique categories that analysis may prove difficult. If such is the case, you might want to utilize the "Storage Used by Workload" analysis described in Section 2.7.2.2, or restrict your analysis to selected groups or classes using either the data selection facilities provided on the primary inquiry screen, or the global data selection exit on the extended options screen.

2.7.2.4 Storage Used by Workload Type-Pie Chart

GRAPHIC RESULTS

This graphic analysis provides a pie chart of total storage frames used by various defined workloads as well as the percentage relative share for each workload.

ANALYSIS GUIDELINES

Total storage includes both central and expanded frames. This chart can help you understand what workloads are your major consumers of storage frames. It can also help you determine if storage usage by a competing workload appears correlated with a period when other measurements reported during that same period were degraded.

You can analyze memory use further using the Real Storage Management analysis reports described in Section 2.2. Virtual storage allocation patterns can have an effect upon processor storage usage. These can be analyzed using the Virtual Storage Management analysis reports found in Section 2.5. Each of these chapters can help you understand and manage memory use by particular workloads.

SPECIAL CONSIDERATIONS

If you do not provide a workload description, as shown in Section 2.7.6.3, each performance group or service class will be treated as a separate workload. This may provide so many unique categories that SAS/GRAPH will shift from using a pie chart to using a horizontal bar chart. While a horizontal bar chart prevents the loss of information from the chart, you will probably find analysis and reporting more meaningful if you describe workloads.

2.7.3 TSO Service Graphs

MVS measures the response time it provides to TSO as well as other workloads. Ordinarily, you would request that MVS manage TSO in a series of performance periods. Generally the size of these periods is based upon the MVS service units consumed. Any given transaction will slip from first period into second after it has consumed a certain number of MVS service units. This continues until the transaction has landed in the last period for the performance group or service class to which it is assigned.

This is a heuristic algorithm--it allows MVS to repeatedly analyze a transaction and learn whether it must continue to treat the transaction as a short, low resource transaction or whether it can begin assigning it priorities for resource usage that your installation has selected for longer running, more resource intensive transactions.

The TSO response graphs allow you to report on both first period TSO response and overall response for the complex of all performance periods.

This section discusses the following graphic analyses:

- 1 - TSO First Period Response
- 2 - TSO Complex Response

2.7.3.1 TSO First Period Response

GRAPHIC RESULTS

This graphic analysis provides a line plot of TSO first period response time.

ANALYSIS GUIDELINES

This plot shows first period TSO response time for each day, week, or month for which data is available.

At the DAYS level of summarization, you should look for consistency. For any day that showed unusual degradation of response time, you should determine whether this was due to a periodic peak load that must be planned for in the future, or whether it was due to some system problem or unusual load that is unlikely to occur frequently.

At higher summarization levels, especially when using the MONTHS level timespan, you can look for trends. If response time is gradually being degraded, you may need to analyze more recent daily TSO response times to look for peaks during the day. You can then use other Performance Manager graphic analyses and reports to search for possible causes for this degradation.

Tracking first period and non-first period TSO response times separately is important because they do not represent the same type of work. First period TSO responses tend to have relatively uniform and small resource requirements, because the SRM or Workload Manager (in the case of systems operating in "goal mode") forces transactions that are heavy resource consumers into subsequent periods. If first period TSO response experiences irregular response times, it is generally due to system load or a system-wide problem.

The overall complex of TSO periods can experience variable response times due to variability of resource usage by some TSO transactions. Tracking overall TSO response time is available with the graph discussed in Section 2.7.3.2.

The detailed behavior of TSO workloads and the system influences that may be affecting them can be further analyzed using the Installation Performance Specification analysis reports and the sysplex analysis reports described in Section 2.8.

SPECIAL CONSIDERATIONS

For systems still utilizing performance groups, you do not need to specify the performance groups that constitute TSO to produce either the plot of TSO first period response time or the plot of overall complex TSO response time. Your TSO performance groups will automatically be chosen for these graphs. However, if you have several performance groups that constitute TSO but have different service level objectives, are they managed differently by the SRM, or do they tend to have markedly different resource requirements? You may use the performance group selection capability found on the main selection screen for either graphic analysis to determine which performance group will be included in a plot.

For systems operating in goal mode, the analysis attempts to determine TSO service classes/periods by searching for the characters "TSO" in either the service class name or the service definition name. If your TSO service classes follow the convention of incorporating TSO in the class name or definition name, the analysis will automatically choose those classes for reporting. Alternatively, you can use the service class selection capability found on the main selection screen of both analyses to determine which service classes will be included in a plot.

2.7.3.2 TSO Complex Response

GRAPHIC RESULTS

This graphic analysis provides a line plot of TSO overall complex response time. This will show you average TSO response times for transactions that may vary considerably in resource requirements.

ANALYSIS GUIDELINES

This plot shows you what sort of variability you are experiencing in overall TSO response times. In conjunction with the TSO first period graph, this plot will help you determine whether consistent first period response times are being maintained in spite of fluctuations in overall complex TSO response times.

TSO non-first period workloads are more closely related to batch type work in one sense: their resource consumption can vary substantially, thus leading to the potential for response time variability. You do not have the kind of control over this response time that you can have over first period response times.

Due to this variability, you may not want to commit to service level objectives for overall TSO response times. Periodically tracking its response tendencies, however, may provide you with early warning about resource constraints or contention issues. If you discover that overall TSO response time is beginning to climb or increase in variability, especially during peak periods, you may want to investigate more fully what the potential causes might be.

The detailed behavior of TSO workloads and the system influences that may be affecting them can be further analyzed using the Installation Performance Specification analysis reports and the sysplex analysis reports described in Section 2.8.

SPECIAL CONSIDERATIONS

For systems still utilizing performance groups, you do not need to specify the performance groups that constitute TSO to produce either the plot of TSO first period response time or the one of overall complex TSO response time. Your TSO performance groups will automatically be chosen for these graphs. However, if you have several performance groups that constitute TSO but have different service level objectives,

are they managed differently by the SRM, or do they tend to have markedly different resource requirements? You may use the performance group selection capability found on the main selection screen for either graphic analysis to determine which performance group will be included in a plot.

For systems operating in "goal mode," the analysis attempts to determine TSO service classes/periods by searching for the characters "TSO" in either the service class name or the service definition name. If your TSO service classes follow the convention of incorporating TSO in the class name or definition name, the analysis will automatically choose those classes for reporting. Alternatively, you can use the service class selection capability found on the main selection screen of both analyses to determine which service classes will be included in a plot.

2.7.4 I/O Tracking Graphs

Analyzing and tracking the performance of disks or direct access storage devices (DASD) is an important part of performance management. The relative speed of moving data to and from DASD is orders of magnitude slower than the speed of moving information in central or expanded storage. Some of the characteristics of DASD make the potential for causing poor system responsiveness even more pronounced.

If two workloads are attempting to access data in central storage, one may have to wait for the other, but once a memory access cycle is made available to the second workload, it will be able to access its data at full memory speed. In contrast, if two workloads are accessing data on a single DASD volume, they may require that the device's read-write heads move back and forth from one file to another. The existence of such heavy seeking activity on a disk can contribute to poor overall system performance or to poor performance of individual workloads.

Even a single TSO or online transaction may cause itself this kind of delay, if it is accessing information at several locations on one device.

This section discusses the following graphic analyses:

- 1 - Device Response Components
- 2 - Average MPL Waiting for Devices

2.7.4.1 Device Response Components

GRAPHIC RESULTS

This graphic analysis provides a stacked vertical bar chart for each of one or more direct access devices that experienced the highest queue times during the interval being charted. The default is to plot only the device showing the highest queue time, but you can choose to analyze more devices per interval reported. The portions of the stacked bar chart represent the average time for the four components of I/O for the device. These components are queue time, pending time, disconnect time, and connect time.

ANALYSIS GUIDELINES

The four components of I/O service time measured by MVS/XA systems are:

- o Queue - The average amount of time that I/O requests to this device were held enqueued within IOS before an SSCH instruction could be issued. Such an enqueue occurs when this particular system already has an I/O operation in progress on a device at the time another I/O operation is requested by this same system.
- o Pending - The time that the I/O was blocked by other I/Os that used some part of the physical path to the device. This can be caused by head of string or control unit contention or reserve delays issued from other systems. Pending time can also include the time that the device was delayed due to an ESCON Director (ESCD) port being busy.
- o Disconnect - The time that the device spent in seeking, in device latency, and in extra rotations due to rotational positioning sensing (RPS) misses.
- o Connect - The time that the device was connected to the path. This includes the time required for both channel control word (CCW) and data transfer.

The first three of these components can be addressed by tuning. Excessive pending time may also indicate a need to adjust the I/O load on the ES Connection Director ports. The fourth component is a function of data set block sizes and can only be addressed by changing the characteristics of the data sets on the device. This last component is the only one that actually represents the movement of data to or from the

device.

The queue component of I/O response time is the sum of the time that the I/Os wait for the device's UCB (that is, device busy condition) and the time that the I/Os wait for the LCU (logical control unit) due to the utilization of the path by other I/Os. If the utilization of the device is more than 30 percent, the I/O queue time can be reduced by spreading the data sets on the volume across two or more devices. If the device utilization is low, the only way to address I/O queuing is to move the volume to a lower activity string.

The pending component of I/O response time represents the average time between successfully issuing the start subcommand (SSCH) to the logical control unit (LCU) and acceptance of the first command by the device. Two factors contribute to this delay. They are the utilization of the path (for example, channels or control units) by other I/Os from the system being measured, or device contention conditions that result from I/Os issued by other systems. These effects may be evaluated using the DASD Skew and Shared DASD Analysis Reports of the Performance Manager.

The disconnect component of I/O response time represents the portion of the device service time during which it is disconnected from the channel. This time includes seek, latency, and extra revolutions that result from RPS missed reconnects. The value shown in the report should be compared to the ideal disconnect time for the device type.

The ideal disconnect time may be calculated by dividing the sum of the average seek time and rotational delay for the device type by two. For example, the average seek time of a 3350 is 25 milliseconds and the rotational delay is 16.7 milliseconds. Therefore, the ideal disconnect time for 3350 type devices is approximately 21 milliseconds.

By comparing the ideal disconnect time to the actual for the device, you can determine the impact of RPS missed reconnects. If this time is significant, you can move the volume to a less active device string where the RPS miss contribution is lower. You can use the DASD Skew Report to evaluate strings where RPS misses are causing performance problems.

Delays due to ESCON director port busy occur when I/O actions are delayed by the temporary unavailability of an ESCON director port. Unless a dedicated port-to-port connection from a channel to a controller has been defined within the

ESCON director, a port-to-port connection will be established dynamically only for the duration of an I/O operation, with the I/O processor requesting termination of the connection upon I/O completion. A channel along which the I/O process is initiated must pass a request that the ESCON director establish a dynamic connection to any of the defined routes to the device for which the I/O operation is intended. Each controller attached to the device is attached to a port on the ESCON director. If each of these ports is busy, an ESCON port busy condition will occur.

If there is direct channel connection to a controller, ESCON port delays are not possible.

It is also possible to link two ESCON directors, for instance, to increase the distance permitted between processors and devices. If this has been done, the link between the directors may also be a source of contention for which the port busy condition will occur.

You can determine the presence of ESCON port delays by using the I/O Component Report of the MVS Performance Manager.

The detailed behavior of I/O devices can be further analyzed using the I/O Configuration analysis reports described in Section 2.4.

SPECIAL CONSIDERATIONS

This graphic analysis uses information that became available with the advent of MVS/XA, so you will not be able to perform this analysis for a system unless it is operating at MVS/XA or beyond.

2.7.4.2 Average MPL Waiting for Devices

GRAPHIC RESULTS

This graphic analysis provides a vertical bar chart for each of one or more direct access devices that experienced the highest queue times during the interval being charted. The default is to plot only the single device showing the highest queue time, but you can choose to analyze more devices per interval reported. The bar chart represents the average number of users of the device that were waiting to gain access to it.

ANALYSIS GUIDELINES

Since this chart shows the average number of applications waiting for a device, you can use it as a warning indicator. The average number of applications waiting is a measurement external to the application. If an application such as one processing online transactions internally queues its own requests for a device, this delay will not show on this chart. This chart will show the queuing that occurs in the I/O subsystem when it has received the I/O request, but is holding it until a previous request that it is managing has completed.

There is no necessarily good or bad number with this chart. You should produce this chart frequently enough to be accustomed to the queuing levels normally seen in your environment.

The detailed behavior of I/O devices can be further analyzed using the I/O Configuration analysis reports described in Section 2.4.

2.7.5 Paging Graphs

When central and expanded storage are sufficiently large, an MVS system may do little or no paging. As workloads grow and applications are enhanced, contention for central and expanded storage may cause MVS to make use of auxiliary storage paging files especially established for this purpose. A limited amount of paging can provide an excellent cushion for central and expanded storage.

However, each time MVS must demand that a page be transferred into central storage from DASD, this represents a dramatic increase in execution time. If the processor attempts to execute an instruction or refers to data that is not in processor storage but has been paged out to auxiliary storage, a processor page fault will occur. MVS will delay the task attempting to execute until the instruction or data it refers to has been moved into central storage. Then MVS will mark the task dispatchable and, depending on the priority of this task relative to other tasks in the system, the task will eventually continue. Thus, the task temporarily runs at the relatively slow speed of DASD I/O instead of at the electronic speed of the processor. The task switching that takes place can also tend to pollute the processor's high speed cache, further degrading performance.

Especially for critical workloads for which you have service level objectives or service level agreements, you need to be assured that demand paging levels are tracked and managed.

This section discusses the following graphic analyses:

- 1 - Response Time for Page Data Set Type
- 2 - Response Time for Local Page Data Sets
- 3 - Demand Page Rates by Performance Group/Svc Class
- 4 - Percent Busy for Local Page Data Sets
- 5 - Demand Page Rates by Workload Type-Bar Chart
- 6 - Demand Page Rates by Workload Type-Pie Chart

2.7.5.1 Response Time for Page Data Set Type

GRAPHIC RESULTS

This graphic analysis provides a multi-line plot of response times for the three types of page data sets: Local, Common, and PLPA (Pageable Link Pack Area).

ANALYSIS GUIDELINES

The response time is the average for all devices being used for a page data set type.

At the DAYS level of summarization, you should look for consistency. For any day that showed unusual degradation of response time, you should determine whether this was due to a periodic peak load that must be planned for in the future, or whether it was due to some system problem or unusual load that is unlikely to occur frequently.

There are essentially only four ways to respond to peak loads and consequent increases in paging response times:

- o Attempt to make a business case for shifting the workload. Although a convincing business case cannot often be successfully made, this can sometimes provide temporary relief for acute paging response problems.
- o Recommend the acquisition of additional memory. This is the most effective and direct, but generally the most costly solution.
- o Add additional paging volumes or improve the performance of the existing ones. This is generally an attractive and not overly costly solution. However, additional paging volumes cannot compensate entirely for a gross overcommitment of processor storage.
- o Attempt to find an offending application and recommend its redesign or other modification. This is often a fruitless endeavor. However, you may occasionally discover an application that has dramatically overcommitted I/O buffers, is searching large arrays in an inefficient manner, or is otherwise misusing processor storage.

At higher summarization levels, especially when using the MONTHS level time-span, you can look for trends. If paging

response time is gradually being degraded, this can give you an early warning that you may need to take some corrective action.

The detailed behavior of the paging subsystem can be further analyzed using the Auxiliary Storage Management analysis reports described in Section 2.2.

2.7.5.2 Response Time for Local Page Data Sets

GRAPHIC RESULTS

This graphic analysis provides a multi-line plot of response times for all local page data sets.

ANALYSIS GUIDELINES

Local page data sets can experience paging loads from a variety of sources: demand paging, swapping, Hiperspace migration, and virtual I/O. Some of this paging traffic is asynchronous, that is, no application is waiting for the page's arrival before it can continue. Some, however, is the result of page faults and will cause the application to have its processing suspended until the page has arrived in processor storage.

With the advent of block paging, which first became available with MVS/ESA SP 4.2, some of this demand paging may lead to a presumptive decision by the Real Storage Manager (RSM) to move into central storage for an address space all pages that meet two criteria: they were previously located in contiguous virtual memory locations and they all had the same unreferenced interval count (UIC) value. This is a heuristic attempt to predict page reference patterns and prepare for them. Initiation of block paging is load dependent, so you may see not only the quantity of pages per unit of time increase during peak loads, but you may also see the nature of the paging work change from short bursts to larger blocks.

The detailed behavior of the paging subsystem can be further analyzed using the Auxiliary Storage Management analysis reports described in Section 2.2.

2.7.5.3 Demand Page Rates by Perf Group/Svc Class

GRAPHIC RESULTS

This graphic analysis provides a stacked vertical bar chart of demand page rates separated by performance group or service class. Each group or class for which there is data is shown as a separate segment on each bar.

ANALYSIS GUIDELINES

If demand paging exists, this chart will show which performance groups or service classes are experiencing the page faults.

When trying to manage demand paging effectively, you should understand what workloads are experiencing it. For MVS systems not yet operating in "goal mode," providing detail by performance group is only effective if workload granularity has been established through two members in the MVS system parameter library, SYS1.PARMLIB. The first of these, IEAIPSxx, the Installation Performance Specifications member, contains the specifications defining to the System Resource Manager of MVS how it is to distribute resources to the various performance groups. The second control member, IEAICSxx, the Installation Control Specifications member, is used by MVS to select the performance group for any particular job, TSO user, started task, or APPC/MVS transaction. Having multiple applications running in the same performance group restricts your capability to distinctly analyze and manage such work.

Similar considerations come into play for MVS environments operating in "goal mode" as introduced in MVS/ESA SP 5.1. Here, the workload granularity must be established through the use of service definitions via the Workload Manager application. Such definitions delineate both service classes, controlling actual resource distribution to different types of work, as well as "report service classes," which can be thought of as similar to the report performance groups defined in IEAICSxx.

Additional information can be obtained by looking at paging response times. The I/O response time multiplied by the number of page faults per second per workload type shows the impact this paging is having on each workload type. Dividing this number by the average MVS multiprogramming level for a workload would give you an estimate of the impact that an individual user or job would experience.

Workloads that offer transaction reporting can provide additional insight into the impact of paging, particularly slow paging, upon these workloads. The portion of the transaction response time caused by paging could be estimated by calculating the number of demand pages per transaction and multiplying the result by the paging response time.

You should understand the behavior of the work in a performance group or service class when demand paging occurs. With a workload like TSO, which places each user in a unique address space, when one user suffers a page fault, the remaining users continue processing, since MVS is managing their address spaces independently. In contrast, when a transaction processing address space, like CICS, experiences a page fault, the impact of MVS halting activity in the address space until the page fault is resolved can be much greater, since the entire MVS task in the CICS address space will go dormant until the page fault is resolved. This has the potential for halting all users that CICS is multitasking at that moment. Only I/O activity in progress for active users will continue.

The detailed behavior of the paging subsystem can be further analyzed using the Auxiliary Storage Management analysis reports described in Section 2.2.

2.7.5.4 Percent Busy for Local Page Data Sets

GRAPHIC RESULTS

This graphic analysis provides a plot of percent device busy for those devices which contain local page data sets.

ANALYSIS GUIDELINES

Local page data sets can experience paging loads from a variety of sources: demand paging, swapping, Hiperspace migration, and virtual I/O. Some of this paging traffic is asynchronous, that is, no application is waiting for the page's arrival before it can continue. Some, however, is the result of page faults and will cause the application to have its processing suspended until the page has arrived in processor storage.

Many people still tend to follow the rule of thumb that says when your page pack reaches 35 percent busy, add another one. For those people, we provide this chart. However, the I/O response time plot discussed in Section 2.7.5.2 provides a better picture of the paging environment, and should be used as the primary guide for analyzing and/or altering the paging configuration.

The detailed behavior of the paging subsystem can be further analyzed using the Auxiliary Storage Management analysis reports described in Section 2.2.

2.7.5.5 Demand Page Rates by Workload Type-Bar Chart

GRAPHIC RESULTS

This graphic analysis provides a stacked vertical bar chart of demand page rates separated by user defined workload types. Each defined workload for which there is data is shown as a separate segment on each bar.

ANALYSIS GUIDELINES

If demand paging exists, this chart will show which workloads are experiencing the page faults.

Additional information can be obtained by looking at paging response times. The I/O response time multiplied by the number of page faults per second per workload type shows the impact this paging is having on each workload type. Dividing this number by the average MVS multiprogramming level for a workload would give you an estimate of the impact that an individual user or job would experience.

Workloads that offer transaction reporting can provide additional insight into the impact of paging, particularly slow paging, upon these workloads. The portion of the transaction response time caused by paging could be estimated by calculating the number of demand pages per transaction and multiplying the result by the paging response time.

You should understand the behavior of the various workload types when demand paging occurs. With a workload like TS0, where each user is placed in a unique address space, when one user suffers a page fault, the remaining users continue processing, since MVS is managing their address spaces independently. In contrast, when a transaction processing address space like CICS experiences a page fault, the impact of MVS halting activity in the address space until the page fault is resolved can be much greater, since the entire MVS task in the CICS address space will go dormant until the page fault is resolved. This has the potential for halting all users that CICS is multitasking at that moment. Only I/O activity in progress for active users will continue.

The detailed behavior of the paging subsystem can be further analyzed using the Auxiliary Storage Management analysis reports described in Section 2.3.

SPECIAL CONSIDERATIONS

If you do not provide a workload description, as shown in Section 2.7.6.3, each performance group or service class will be treated as a separate workload. Further, if you choose to use report performance groups (rather than control groups), or report service classes (rather than standard classes), overlapping resource utilization of such groups or classes will tend to overstate actual utilization, and should only be done with a full understanding of the impact on the resulting analysis.

2.7.5.6 Demand Page Rates by Workload Type-Pie Chart

GRAPHIC RESULTS

This graphic analysis provides a pie chart of demand page rates and relative shares belonging to user defined workload types. Each defined workload for which there is data is shown as a separate slice.

ANALYSIS GUIDELINES

If demand paging exists, this chart will show which workloads are experiencing the page faults.

Additional information can be obtained by looking at paging response times. The I/O response time multiplied by the number of page faults per second per workload type show the impact this paging is having on each workload type. Dividing this number by the average MVS multiprogramming level for a workload would give you an estimate of the effect that an individual user or job would experience.

Workloads that offer transaction reporting can provide additional insight into the impact of paging, particularly slow paging, upon these workloads. The portion of the transaction response time caused by paging could be estimated by calculating the number of demand pages per transaction and multiplying the result by the paging response time.

The detailed behavior of the paging subsystem can be further analyzed using the Auxiliary Storage Management analysis reports described in Section 2.3.

SPECIAL CONSIDERATIONS

If you do not provide a workload description, as shown in Section 2.7.6.3, each performance group or service class will be treated as a separate workload. This may provide so many unique categories that SAS/GRAPH will shift from using a pie chart to using a horizontal bar chart. While a horizontal bar chart prevents the loss of information from the chart, you will probably find analysis and reporting more meaningful if you describe workloads.

A further consideration is that if you choose to use report performance groups (rather than control groups), or report service classes (rather than standard classes), overlapping resource utilization of such groups or classes will tend to overstate actual utilization, and should only be done with a

full understanding of the impact on the resulting analysis.

2.7.6 Producing Graphs

All the graphic analyses are produced by using the online menus and selection panels of the CA MICS Performance Manager Option available from within the CA MICS Workstation Facility (MWF). Chapter 1 of this guide assists you in accessing the online menus and selection panels if you are not familiar with MWF. The following sections will discuss the use of graphic selection panels to produce the performance analysis and tracking graphics you need.

This section discusses the following topics:

- 1 - Overall Operation
- 2 - Optional Selection Criteria
- 3 - Workload Groupings
- 4 - Color Chart Options
- 5 - Alternate Input Files
- 6 - Saved Output Files

2.7.6.1 Overall Operation

Making no changes to a graphic analysis selection screen prior to executing the inquiry will cause default selections to be used to produce a graph. If no changes are made, the result will be produced as a line printer graph. You may find that producing one of these character-based graphs is a good starting point for determining what, if any, options you need to specify to get the kind of graph that you will find most helpful.

Figure 2-200 shows a typical main selection screen that you would use to specify various choices prior to producing a graph. If you need help using the Performance Manager Option, see Section 1.4 of this guide. The selection options present for all types of MVS Performance Manager graphic analyses are:

- o Timespan

Use this option to select which CA MICS timespan will supply the data for your graph. Setting this to a higher timespan (WEEKS or MONTHS) can provide you with a longer-range overview, whereas setting it to the DETAIL or DAYS timespan can help you focus on a particular problem that needs detailed analysis.

- o Cycles

Select the particular cycle or cycles of data that will be used to produce your graph. Cycle 01 of the DETAIL or DAYS timespan would represent the information collected by during its most recent daily update processing. Cycle 01 of the WEEKS or MONTHS timespan would represent the information summarized for the previous week or month. You may specify a single cycle or a range of cycles.

- o Allow Run-Time Override

Both the timespan and the cycles can be entered on this selection screen and used directly to specify processing options. Alternatively, you can choose to allow the run-time override of either or both of these selections. You may find this convenient if you produce a particular graph on an adhoc basis. The same ability to allow run-time overrides exist for many of the other selection options available for graphic analysis.

- o Report by Entire Sysplex

Some of the CPU activity graphs give you the option of summarizing activity across all of the systems that comprise a Parallel Sysplex. This option requires that the CA MICS input files contain the data element SYSPLEX. The CA MICS Hardware and SCP Analyzer product must be generated with the MVS 5.1 Coupling Facility Cross-System Extended Service data elements activated. See the CA MICS Hardware and SCP Analyzer guide for more information on activating these elements.

- o Chart by hour if timespan is DAYS or WEEKS

Many of the graphs permit you to select whether to plot individual hours as separate data points, or to summarize multiple adjacent hours into a single summarized point.

- o Generate a graph for each value of zone

You can choose to produce a separate graph for each CA MICS zone for which there is data or you can choose to combine multiple zones into a single graph.

- o Generate color graphic pie charts

You can specify that a color graphic plot be produced on a graphic device available to MICF or you can permit the graph to be produced as a character-based graph on a line printer. If you are unfamiliar with producing graphic results using MICF, you may want to contact your CA MICS administrator to determine if any special MICF options are needed.

- o Graph subtitle

You can optionally specify a subtitle that will be included on the graph.

- o Type of performance group or service class

MVS has two kinds of performance groups: control (CTL) and report (RPG) performance groups. With the introduction of MVS/ESA SP 5.1, systems operating in what is known as "Goal Mode" have two kinds of service classes: standard (STD) and report (RSC) service

classes.

For pre-SP 5.1 systems, and for SP 5.1 and above systems continuing to operate in "compatibility mode," the classification of jobs, TSO users, started tasks, and APPC/MVS transactions into performance groups is ordinarily done in the SYS1.PARMLIB(IEAICSxx) system library member. If a graphic analysis is based on performance group information, you need to specify whether you want to use information gathered for control performance groups or for report performance groups. This selection option appears for many, but not all, of the graphs, since some of them are not based on MVS performance group information.

For MVS/ESA SP 5.1 and systems operating in goal mode, the Workload Manager is used to classify jobs, TSO users, etc., into service classes via the specification of service definitions. Standard service class definitions control the distribution of resources to selected workload types. Report service class definitions do not have an effect on Workload Manager's distribution of resources, but can be used to capture resource and service measures for specific types of work.

Several data selection criteria are available for graphic analyses. If any of these selection criteria are used, they will serve to limit the information used in the graphic analysis:

- o Date

Two date ranges are available. In each range either a single date or a pair of dates can be entered using a DDMMYY format, where MON is the first three letters of the month.

- o Hour

Three hour ranges are available. In each range either a single hour or a pair of hours can be entered using values from 00 to 23.

- o Zone

Nine zone selection options are available. It is common, but not universal, for zone 1 to be defined to include prime shift hours for weekdays. If this were

true for your CA MICS system, entering 1 into zone selection would limit graphic analysis to data from this period. Your CA MICS administrator can assist you if you are unfamiliar with how zones have been defined in your CA MICS system.

- o SYSID

Six SYSID selection options are available. You can limit your graphic analysis to certain SYSIDs by making use of these selection options. Your CA MICS administrator can assist you if you are unfamiliar with how SYSIDs have been defined in your CA MICS system.

- o Performance Group

This selection option appears for many, but not all, of the graphs, since some of them are not based on MVS performance group information. If it is displayed, eight performance group selection options are available. You can limit your graphic analysis to certain performance groups by making use of these selection options. The definition of performance groups is ordinarily done in the IEAICSxx member of the SYS1.PARMLIB system library. The person responsible for maintaining this system library member may be able to assist you if you are unfamiliar with how performance groups have been defined in your MVS systems.

- o Service Class

Any analysis that can have performance group selection can also have service class selection to achieve the same sort of restricted analysis. If it is displayed, four service class selection options are available. The definition of service classes is done via the Workload Manager application. The person responsible for maintaining Workload Manager may be able to assist you if you are unfamiliar with how service classes have been defined in your MVS systems.

Note that you cannot use both Performance Group and Service Class selection in the same analysis inquiry. The "Type of performance group or service class" option dictates which of these two data selection criteria is available. The practical implication of this is that if you have systems operating in both compatibility mode and in goal mode,

separate analyses will be needed to address the two different environments.

```

/----- CPU Busy by Workload Type -----\
|Command ==>
|
|Modifying CA MICS Inquiry: CPU001 - CPU Analysis of Workloads
|
|Time-span ==> DAYS (DETAIL/DAYS/WEEKS/MONTHS) Cycle(s) ==> 01 - 05
|Data base ==> P (PRODCPLX) _ _ _ _ _
|Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N)
|Report by Entire Sysplex ==> N (Y/N)
|Chart by hour if timespan is DAYS or WEEKS ==> N (Y/N)
|Generate a pie chart for each value of Zone ==> Y (Y/N)
|Generate color graphic pie charts ==> N (Y/N)
|Graph subtitle ==> This Week's Utilization Patterns
|Type of performance group or service class ==> CTL (CTL, RPG, STD, or RSC)
|
|Data Selection: Allow run-time override
|Date ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ ==> N (Y/N)
|Hour ==> 07 - 20 _ _ - _ _ _ _ - _ _ ==> N (Y/N)
|Zone ==> 1 _ _ _ _ _ ==> N (Y/N)
|Sysid ==> S002 _ _ _ _ _ ==> N (Y/N)
|Perf. Group ==> _ _ _ _ _ ==>
|Service Class ==> _ _ _ _ _
|Specify extended options ==> N (Y/N)
\-----/

```

Figure 2-200. Graphic Inquiry Main Selection Screen Example

o Specifying extended options

Specifying "Y" for extended options on the main graphic selection screen shown in Figure 2-200 causes the display of a menu (shown in Figure 2-201) that permits you to choose a variety of additional options for data manipulation or color graphic processing:

- Specify statements for global data selection
- Specify statements for workload definition
- Specify alternate input file(s)
- Save the selected input observations
- Override color graphics format parameters
- Choose the color graphics format set

Each of these extended options is discussed in sections that follow, except for the one that deals


```

|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy
|
|Cmd  SAS Statements for: Specifying global data selection
| - -----
|_ - %MACRO SELECT;
|
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
|_ - %MEND SELECT;
|***** BOTTOM OF DATA *****
|
|-----/

```

Figure 2-202. Global Data Selection Exit Screen Example

- o Specifying a global data selection exit

The first item on the extended option menu shown in Figure 2-201 permits you to specify data selection using one or more lines of SAS code. If you enter "Y" (Yes) for the Specify global data selection exit selection in the Extended Options Menu, the Global Data Selection Exit screen shown in Figure 2-201 is displayed. Any valid SAS statements can be used here to cause the deletion of observations you do not want included in any particular graphic analysis.

Ordinarily, the standard selection options available on the main graphic selection screen for any particular graph should provide sufficient selection capability. However, this exit is available for any special circumstances not covered by these standard selection options.

2.7.6.2 Optional Selection Criteria

As mentioned in the previous section, performance group/service class selection options are common but not universal among the graphic analyses. Several other selection options mentioned in the previous section may not be present on the main selection screen for any particular graph if they are not appropriate for that graph.

Graphs with "Workload" in their title permit you to define workload groupings. This capability is described in the next section.

Certain of the graphs depend upon information that is available in the form required only at the DETAIL or DAYS timespans. The valid CA MICS timespans for any graph are shown on its main selection screen.

2.7.6.3 Workload Groupings

For those graphs that contain the word "Workload" in their title, two alternative extended options allow you to define workload groupings. You can specify a workload definition exit, which will permit you to write SAS statements combining homogeneous performance groups or service classes into unique workload groups.

If you enter Y (Yes) for the Specify workload definition exit selection in the Extended Options Menu, the Workload Grouping Exit screen shown in Figure 2-203 is displayed.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> CSR |
|
|Inquiry Step: CPU Busy by Workload Type
|
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy
|
|Cmd  SAS Statements for: Specifying workload grouping exit
| -  -----
| -  %MACRO GROUP;
|    /* SAMPLE GROUP EXIT LOGIC */
|    IF PERFGROUP = 2 THEN GROUP = 'TSO';
|    ELSE IF PERFGROUP = 1 OR PERFGROUP = 23 THEN GROUP = 'BATCH';
|    ELSE IF PERFGROUP = 49 THEN GROUP = 'CICS';
|    ELSE GROUP = 'OTHER';
|    -----
|    -----
|    -----
|    -----
|    -----
| -  %MEND GROUP;
|***** BOTTOM OF DATA *****|
|
\-----/

```

Figure 2-203. Workload Grouping Exit Screen Example

Specifying a workload grouping exit is an option for any graphs whose main selection screen mentions workload types. The screen shown in Figure 2-203 permits you to specify SAS code to define values for the SAS data element named GROUP. This is the display you would see if the data source for the analysis is performance group data for a pre-MVS/ESA SP 5.1 system, or an SP 5.1 or above system operating in

compatibility mode. If you do not provide a workload grouping exit for any of the graphic analyses that make use of this exit, GROUP will be set equal to the value of PERFGRP, the performance group. Using this exit permits you to collapse several performance groups into one workload group, thus providing a convenient method for summarizing related workloads together.

For analyses which use service class data from MVS/ESA SP 5.1 and above systems operating in goal mode, a slightly different screen will appear. The %MACRO is named CLASS, and the values you would define would be for the data element CLASS. So, for example, the following workload grouping exit code might group specific standard service classes into reporting workloads:

```
%MACRO CLASS;
  SELECT (WLMCLASS);
    WHEN ('SYSTEM') CLASS='SYSTEM';
    WHEN ('ASCH') CLASS='SYSTEM';
    WHEN ('SYSSTC') CLASS='SYSTEM';
    WHEN ('MVSSUBSY') CLASS= 'SYSTEM';
    WHEN ('CICSREGS') CLASS= 'CICS';
    WHEN ('IMSREGS') CLASS='IMS';
    WHEN ('BATCH') CLASS='BATCH';
    OTHERWISE CLASS = 'OTHER';
  END;
%MEND CLASS;
```

If the data source for the analysis is report service classes, rather than standard service classes, then element WLMRPTCL would be used instead of WLMCLASS. Also, if you do not provide a workload grouping exit, CLASS will be set equal to the value of either WLMCLASS or WLMRPTCL depending on the data source (Standard or Report) specified on the primary screen.

2.7.6.4 Color Chart Options

Each graphic analysis of the MVS Performance Manager creates one of three types of graphs: a line plot, a vertical bar chart, or a pie chart. Each of these types of graphs has slightly different options available.

This section discusses the following topics:

- 1 - Color Graphic Device Selection
- 2 - Line Plot Options
- 3 - Bar Chart Options
- 4 - Pie Chart Options

2.7.6.4.1 Color Graphic Device Selection

COLOR GRAPHICS FORMAT SET

This parameter is required when you specify "Y" for the color graphics format parameters extended option. It is used to fill in a set of color and other choices for graphic plotting. If you have not specified the use of the color graphics format parameters extended option, this parameter is ignored.

Multiple sets of color graphics format parameters, called color graphics format sets, are supported by MICF. You will probably want to have a unique format set for each different graphic device you use, because each one might display particular colors and fonts most favorably. When you invoke the color graphics format parameters extended option, parameter defaults are taken from the specified color graphics format set. Changing the format set specified will cause the color graphics format parameters extended option to be reset to the new defaults.

This parameter is set by default to the color graphics format set specified for the MICF default execution mode (e.g., either the default for foreground or the default for batch execution.)

You may specify ? (question mark) to display a selection list of available color graphics format sets. The Graphics Format Set Selection panel is displayed when you invoke the color graphics format parameters extended option.

2.7.6.4.2 Line Plot Options

LINE PLOT PARAMETERS

You can use the color graphics format parameters extended option to override parameter defaults that control color graphics format and appearance. These default values are specified through the Color Graphics Format Parameters process of MIF Options. Thus, you can customize these defaults to match your most common reporting requirements. In this way, you need to override color graphics format defaults only for exceptional reporting requirements.

The Color Graphics Format Parameters panels for line plots include specifications for the following:

- o Colors for axis, text, title, bars, background, and other plot format related parameters.
- o Color, symbol, and line type for each plot line.

An example of a Color Graphics Format Parameter panel for line plots can be seen in Figure 2-205. Its parameters are described below.

```

/----- Color Graphics Format Parameters -----\
|Command ==>                                     Format set: IBM3825 |
|
|Inquiry step: CPU Busy by Hour and Day
|
|Title print font      ==> TRIPLEX      Title color      ==> BLACK
|Subtitle Print font   ==> SIMPLEX      Subtitle color    ==> BLACK
|Draw border around plot ==> N (Y/N) Background color    ==> -----
|Axis color            ==> BLACK        Area inside axis color ==> -----
|Text color            ==> BLACK
|
|1st Day:              Color ==> RED      Sym ==> NONE      Line ==> SOLID
|2nd Day:              Color ==> BLUE     Sym ==> NONE      Line ==> DASHED
|3rd Day:              Color ==> BLACK    Sym ==> NONE      Line ==> DOTTED
|4th Day:              Color ==> GREEN    Sym ==> NONE      Line ==> SOLID
|5th Day:              Color ==> RED      Sym ==> NONE      Line ==> SOLID
|6th Day:              Color ==> BLUE     Sym ==> NONE      Line ==> SOLID
|7th Day:              Color ==> BLACK    Sym ==> NONE      Line ==> SOLID
|8th Day:              Color ==> GREEN    Sym ==> NONE      Line ==> DOTTED
|9th Day:              Color ==> RED      Sym ==> NONE      Line ==> DASHED
|
|Symbol(NONE,PLUS,STAR,TRIANGLE,DIAMOND,HASH,SQUARE,A-Z,0-9,or SAS/GRAPH value)
|Color (May be any color available on your graphics device.)
|Line (SOLID,DASHED,DOTTED,1-32)
\-----/

```

Figure 2-205. Plot Graphics Format Parameters Panel Example

GENERAL LINE PLOT PARAMETER CHOICES

- o Print Fonts - These parameters specify the character format (or font) for titles and subheadings. Character font names are 1-9 characters. Refer to the SAS/GRAPH User's Guide for a list of valid character fonts.

- Title - Required.

- This refers to the primary title line.

- Subtitles - Required.

- This refers to the subtitles and subheading line.

- o Draw border - Required.

- Indicates whether a border should be drawn around the plot. Default is N (No).

- o Colors

- This group of parameters specifies the color to use for portions of the plot. Each parameter is an 8-character color name. This parameter is not validated; however, specifying an invalid color will cause execution to fail. Refer to the SAS/GRAPH User's Guide for valid colors for each graphics device.

- Title - Required.

- This refers to the primary title line.

- Subtitles - Required.

- This refers to the subtitles and subheading line.

- Axis - Required.

- This refers to the plot axes (X-axis and Y-axis).

- Text - Required.

This applies to all text (words) except the titles and subheading (e.g., axis labels, symbol legend, etc.)

- Background - Optional.

Specify this parameter to generate the graph on a colored background (fill the entire graph with a color).

- Area Inside the Axis - Optional.

Specify this parameter to draw the plots on a colored background. This is similar to the background parameter, except only the area inside the axes is filled.

LINE PLOT PARAMETER CHOICES

The following parameters define the color, symbol, and line type that will appear on the plot. The graphics parameters are active in all plots. The three basic parameters are described below.

- o Color - character color name(s).

This parameter is not validated; however, specifying an invalid color will cause execution to fail. Refer to the SAS/GRAPH User's Guide for valid colors for each graphics device.

- o Symbol - character symbol value(s).

Specify NONE to have no symbol depicting the data points on the plot. Specify PLUS, STAR, TRIANGLE, DIAMOND, HASH, SQUARE, letters A-Z, or digits 0-9 to use as the plot character for data points on the plot. See the SAS/GRAPH User's Guide for additional symbol specifications.

- o Line - character line type(s).

Specify SOLID, DOTTED, DASHED, or a number from 1-32 indicating the type of line that will be drawn between data points on the plot. The SAS/GRAPH User's Guide shows the different kinds of lines that can be produced when numbers 1-32 are used.

2.7.6.4.3 Bar Chart Options

BAR CHART PARAMETERS

You can use the color graphics format parameters extended option to override parameter defaults that control color graphics format and appearance. These default values are specified through the Color Graphics Format Parameters process of MIFC Options. Thus, you can customize these defaults to match your most common reporting requirements. In this way, you need to override color graphics format defaults only for exceptional reporting requirements.

The Color Graphics Format Parameters screen includes specifications for the following:

- o Colors for axis, text, title, bars, background, and other chart format related parameters.
- o Color and Pattern for portions of each bar representing CPU utilization.

An example of a Color Graphics Format Parameter screen for bar charts can be seen in Figure 2-206. Its parameters are described below.

```

/----- Color Graphics Format Parameters -----\
|Command ==>                                     Format set: IBM3825 |
|
|Inquiry step: CPU Busy by Performance Group      |
|
|Title print font      ==> XSWISSB      Title color      ==> BLACK |
|Subtitle Print font   ==> NONE         Subtitle color    ==> BLACK |
|Draw border around plot ==> N (Y/N)   Print Symbol Legend ==> Y (Y/N) |
|Axis color            ==> BLACK        Area inside axis color ==> ----- |
|Text color            ==> BLACK        Background color   ==> ----- |
|
|Perf. Group
|by resources  Color   Pattern
|-----
| 1st         RED     SOLID      Color may be any color available on your
| 2nd         GREEN   RIGHTHASH  graphics device.
| 3rd         BLUE    LEFTHASH
| 4th         BLACK   HASHED     Pattern may be SOLID, HASHED, RIGHTHASH,
| 5th         RED     SOLID      LEFTHASH, HOLLOW, or any valid SAS/GRAPH
| 6th         GREEN   RIGHTHASH  pattern value.
| 7th         BLUE    LEFTHASH
| 8th         BLACK   HASHED
| Others      RED     HOLLOW
\-----/

```

Figure 2-206. Bar Chart Color Graphics Format Parameters Screen Example

GENERAL BAR CHART PARAMETER CHOICES

- o Print Fonts - These parameters specify the character format (or font) for titles and subheadings. Character font names are 1-9 characters. Refer to the SAS/GRAPH User's Guide for a list of valid character fonts.

- Title - Required. This refers to the primary title line.

- Subtitles - Required. This refers to the subtitles and subheading line.

- o Draw border - Required.

Indicates whether a border should be drawn around the bar chart. Default is N (No).

- o Print Legend - Required.

Allows you to include a legend in the chart indicating which performance groups or service classes are represented by each bar segment.

- o Colors

This group of parameters specifies the color to use for portions of the chart. Each parameter is an 8-character color name. This parameter is not validated; however, specifying an invalid color will cause execution to fail. Refer to the SAS/GRAPH User's Guide for valid colors for each graphics device.

- Title - Required.

This refers to the primary title line.

- Subtitles - Required.

This refers to the subtitles and subheading line.

- Axis - Required.

This refers to the chart axes (X-axis and Y-axis).

- Text - Required.

This applies to all text (words) except the titles and subheading (for example, axis labels, symbol legend).

- Background - Optional.

Specify this parameter to generate the graph on a colored background (fill the entire graph with a color).

- Area Inside the Axis - Optional.

Specify this parameter to draw the bars on a colored background. This is similar to the background parameter, except only the area inside the axes is filled.

LINE PLOT PARAMETER CHOICES

The following parameters define the color and pattern for each bar appearing in the chart. The graphics parameters are active in all charts. The two basic parameters are described below.

- o Color - character color name(s).

This parameter is not validated; however, specifying an invalid color will cause execution to fail. Refer to the SAS/GRAPH User's Guide for valid colors for each graphics device.

- o Pattern - character pattern value(s).

Specify SOLID to fill the area with the corresponding color. Specify HOLLOW to simply outline the area in the corresponding color. Specify HASHED to fill the area with lines in a cross-hatched pattern using the corresponding color (RIGHTHASH and LEFTHASH are each half of the HASHED pattern.) See the SAS/GRAPH User's Guide for additional pattern specifications.

2.7.6.4.4 Pie Chart Options

PIE CHART PARAMETERS

You can use the color graphics format parameters extended option to override parameter defaults that control color graphics format and appearance. These default values are specified through the Color Graphics Format Parameters process of MIF Options. Thus, you can customize these defaults to match your most common reporting requirements. In this way, you need to override color graphics format defaults only for exceptional reporting requirements.

The Color Graphics Format Parameters panel includes specifications for the following:

- o Colors for text, title, background, and other chart format related parameters.
- o Color and Pattern for portions of each pie slice representing a given workload's share of CPU consumption.

An example of a Color Graphics Format Parameter panel for pie charts can be seen in Figure 2-207. Its parameters are described below.

```

/----- Color Graphics Format Parameters -----\
|Command ==>                                     Format set: IBM3825 |
|
|Inquiry step: CPU Busy by Workload Type
|
|Title print font      ==> TRIPLEX      Title color      ==> BLACK
|Subtitle print font   ==> SIMPLEX     Subtitle color    ==> BLACK
|Draw border around plot ==> N      (Y/N)
|Background color      ==> _____ Text color        ==> BLACK
|
|Wkl. Group
|by resources  Color    Pattern
|-----
| 1st          _____ SOLID      Color may be any color available on your
| 2nd          _____ RIGHTHASH  graphics device.
| 3rd          _____ LEFTHASH
| 4th          _____ HASHED     Pattern may be SOLID, HASHED, RIGHTHASH,
| 5th          _____ SOLID      LEFTHASH, HOLLOW, or any valid SAS/GRAPH
| 6th          _____ RIGHTHASH  pattern value.
| 7th          _____ LEFTHASH
| 8th          _____ HASHED
| Others       _____ SOLID
|

```

\-----/

Figure 2-207. Pie Chart Color Graphic Format Parameters Panel Example

GENERAL PIE CHART PARAMETER CHOICES

- o Print Fonts - These parameters specify the character format (or font) for titles and subheadings. Character font names are 1-9 characters. Refer to the SAS/GRAPH User's Guide for a list of valid character fonts.

- Title - Required.

- This refers to the primary title line.

- Subtitles - Required.

- This refers to the subtitles and subheading line.

- o Draw border - Required.

Indicates whether a border should be drawn around the pie chart. Default is N (No).

- o Colors

This group of parameters specifies the color to use for portions of the chart. Each parameter is an 8-character color name. This parameter is not validated; however, specifying an invalid color will cause execution to fail. Refer to the SAS/GRAPH User's Guide for valid colors for each graphics device.

- Title - Required.

- This refers to the primary title line.

- Subtitles - Required.

- This refers to the subtitles and subheading line.

- Text - Required.

- This applies to all text (words) except the titles and subheading (e.g., slice labels, symbol legend,

etc.)

- Background - Optional.

Specify this parameter to generate the graph on a colored background (fill the entire graph with a color.)

PIE SLICE PARAMETER CHOICES

The following parameters define the color and pattern for each slice appearing in the pie chart. The graphics parameters are active in all charts. The two basic parameters are described below.

- o Color - character color name(s).

This parameter is not validated; however, specifying an invalid color will cause execution to fail. Refer to the SAS/GRAPH User's Guide for valid colors for each graphics device.

- o Pattern - character pattern values(s).

Specify SOLID to fill the area with the corresponding color. Specify HOLLOW to simply outline the area in the corresponding color. Specify HASHED to fill the area with lines in a cross-hatched pattern using the corresponding color (RIGHTHASH and LEFTHASH are each half of the HASHED pattern). See the SAS/GRAPH User's Guide for additional pattern specifications.

2.7.6.5 Alternate Input Files

If you enter "Y" (Yes) for the Specify alternate input file selection in the Extended Options Menu, the Alternate Input File panel shown in Figure 2-208 is displayed.

```
/----- Alternate Input File -----\
|Command ==>
|
|Inquiry Step: CPU Busy by Workload Type
|
|Specified Time-span: DAYS
|
|Equivalent CA MICS file(s): SCPPGA HARCPU
|Alternate SAS file(s) ==> _____
|
|SAS data set allocation instructions:
|Data set name ==> _____
|Disposition ==> SHR (SHR/OLD)
```

Figure 2-208. Example of Alternate Input File Panel

This panel allows you to specify a SAS file that is not a part of the standard CA MICS online database as input to the graphic analyses. You must specify one or more SAS file names and an MVS data set name. If you do not enclose the data set name in single quotes, your TSO prefix (or user-ID) is added as the first portion of the data set name.

The number of alternative SAS files required and the equivalent CA MICS file names will automatically be adjusted for the analysis you are selecting.

2.7.6.6 Saved Output Files

If you enter "Y" (Yes) for the Save selected input observations selection in the Extended Options Menu, the Save selected input observations panel shown in Figure 2-209 is displayed.

```
/----- Save selected input observations -----  
|Command ==>  
|  
|Inquiry Step: CPU Busy by Workload Type  
|  
|Equivalent CA MICS file(s): SCPPGA HARCPU  
|Saved SAS file name(s) ==> _____  
|  
|SAS data set name      ==> _____  
|  
|  
|  
|  
|  
|  
|
```

Figure 2-209. Save Selected Input Observations Panel Example

This panel allows you to specify a SAS file that is not a part of the CA MICS online database into which output from the graphic analyses will be saved. Use this option if you want to save a copy of the data that was used to produce the graph.

You must enter the SAS file name(s) and an MVS data set (SAS data set name) that will be used to contain the SAS files. During a subsequent execution of this graphic analysis, you can read this same data for processing by using the Specify alternate input file extended option as described previously.

If you do not enclose the data set name in single quotes,

your TSO prefix (or user-ID) is added as the first portion of the data set name.

The number of alternative SAS files required and the equivalent CA MICS file names will automatically be adjusted for the analysis you are selecting.

2.8 Sysplex Analysis

The sysplex analysis (SPA) routines present reports from the RMF hardware and system control program files of the CA MICS database to provide a convenient method of tracking the operation of a sysplex and the underlying coupling facility linking the various systems together.

IBM's System/390 Parallel Sysplex consists of hardware (the Coupling Facility) and software (starting with MVS/ESA SP 5.1) combining to provide high performance data sharing and a single point of control of all systems configured in the sysplex.

The S/390 Coupling Facility provides services necessary for multisystem data sharing by authorized systems in a parallel sysplex. It does this by handling serialization and buffer management, caching for shared data, and message queuing. The result is the appearance of a single system and a single database in a parallel group.

The modifications to MVS in support of a parallel sysplex include enhancements that provide more simplified management of multiple systems, and new functions to support additional open standards interfaces. Beyond these changes, many of the MVS subsystems have also been modified to take advantage of data sharing and system sharing services, including CICS, IMS, JES, and OpenEdition/MVS.

Needless to say, the simplified outward appearance of a sysplex environment is accomplished through very sophisticated hardware and software operating together to insure data integrity and reliability. Because multiple systems are operating conjunctively, considerable overhead must be dedicated to communication between systems, and the handling of virtually any situation that might interfere with the smooth and efficient processing demanded by today's information systems organizations.

The goal of the sysplex analysis software is to ensure that the parallel sysplex is optimally configured, and that adequate resources have been allocated across the entire sysplex. The reports allow you to map the existing configuration; determine what messaging activity is taking place between and within individual systems; and to detect unusual or out of the ordinary conditions that may be impacting on over-all system throughput and availability, including suggestions on the course of action that should be taken.

Section 2.8.4 defines the terms and explains the basic dynamics of many of the elements used in sysplex analysis. Depending on your familiarity with these terms and concepts, you may find it useful to read this section before proceeding with Sections 2.8.2 and 2.8.3.

The following sections describe the operation of the sysplex analysis routines and how to use the reports they produce:

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output
- 4 - Parallel Sysplex Concepts
- 5 - Component Operation

2.8.1 Functional Description

The sysplex analysis portion of the CA MICS Performance Manager Option provides both software and a methodology to evaluate the operation of a parallel sysplex. The software consists of reports and tables to measure the interactivity between members, resource utilization on the Coupling Facility, and potential performance problems and bottlenecks.

Explanations of the terms used below in describing these reports and tables are presented in Section 2.8.4. A description of the controls available and how to evaluate the effect of those controls on the various components of the sysplex is given in the standard output and the parallel sysplex concepts sections (see Sections 2.8.3 and 2.8.4).

The range of data that can be reported from a particular CA MICS database depends on the timespans that were defined for these files and the number of cycles defined in each timespan. If you are not familiar with the generation options for the database unit containing the MVS SCP Information Area, ask your CA MICS System Administrator for this information.

Five basic reports are available. Each report can be tailored during online definition to display data at various timespans and summarization patterns (reporting by zone, hour, or day, etc.). The reports produced are:

- o Sysplex Configuration Map - Reports that XCF members are signaling other members. Both total signals and signaling raters are reported, as well as indicators of externally (operations) induced alterations to the sysplex environment such as member activations and deactivations, and resetting of the counters maintained by the system during the interval covered by the report. The report is designed to aid in understanding how the sysplex is structured, and to help in planning for an optimal configuration.
- o Sysplex XCF Member Profile - Provides a detailed report on the activity for a specific XCF Member. The report shows, by the selected time interval, both environmental factors guiding member behavior (activations, deactivations, etc.), and utilization measures addressing signaling activity both for the selected member and the whole sysplex. The members relative contribution to overall signaling activity is also presented. The report helps with determining activity

patterns within specific members, and how those patterns may be contributing to overall sysplex performance.

- o Sysplex Performance Exceptions Report - Reports on unusual or out of the ordinary conditions that may represent a performance problem for the sysplex. Information reported includes the interval and duration during which the exception(s) was/were detected, the number of occurrences and an explanation or recommended course of action for addressing the exception.
- o Sysplex XCF Channel Activity Report - Displays the channel activity associated with the Coupling Facility. This report shows, by the selected time interval, key measures of request activity on the channel subsystem, including rate of activity, number of unsuccessful requests, and average service time. Additionally, average path and subchannel availability and usage is reported. The report can be used to track subsystem resource utilization, and can help with determining when increased capacity may be needed.
- o Coupling Facility Configuration Report - This report provides a table depicting interval by interval Coupling Facility CPU activity and configuration changes. Average and maximum percent utilization of the Coupling Facility is reported, as is the amount of storage allocated and used by the defined structures. The report provides feedback to help with tuning the Coupling Facility.

2.8.2 Usage Guidelines

This section helps you use the sysplex analysis functions of the CA MICS Performance Manager. The following topics are discussed:

- 1 - Tracking Sysplex XCF Services
- 2 - Tracking Coupling Facility Hardware
- 3 - Exception Review and Management

2.8.2.1 Tracking Sysplex XCF Services

Fundamental to making decisions about member groups and transport class assignment is the need to know what members are signaling each other, what volume of signaling activity is taking place, and what are the physical characteristics of the signaling activity.

The performance of the systems in the sysplex is closely related to the performance of the signaling service. Throughput and response of participating sysplex system is quite dependent on the signaling service.

The major items that need to be tracked in the signaling environment are:

- o What members are signaling other members
- o The appropriateness of transport classes in terms of XCF groups assigned and message (buffer) lengths being processed
- o Whether the number of signaling paths provided is appropriate for the amount of message traffic.

The sysplex configuration map and the XCF member profile reports are designed to help in assessing these important issues. Use the configuration map to assess the first item above, and then utilize the member profiles to examine more closely the activity and performance experienced by the individual XCF members.

2.8.2.2 Tracking Coupling Facility Hardware

It should come as no surprise that as more MVS systems participate in the data sharing activities of a parallel sysplex, the greater is the need for increased coupling facility capacity and coupling facility channel connectivity.

As you install or add coupling facility channels between MVS systems and coupling facilities, you must consider the performance and availability ramifications for both the hardware and software. On the hardware side, there are issues such as the amount of memory to include in the coupling facility, how it's allocated among various structures, and the number of coupling facility channels for each central processor complex (CPC). On the software side, issues revolve around the number of software products that will use the coupling facility channels, and the kind of data sharing you plan to implement in the sysplex.

As you gradually build up your sysplex, it is advisable to incrementally add applications and workload to the data sharing environment. As you do so, monitor the structure sizes, request types (synchronous/asynchronous) and volume, queuing levels, and service times.

The XCF Channel Activity report can be used to assess overall request activity on the sysplex, as well as path and subchannel availability. The Coupling Facility Configuration report can be used to more closely examine the request activity going through particular coupling facility list structures.

2.8.2.3 Exception Review and Management

The exception reporting capabilities provided with the sysplex analysis function are designed to provide a concise, integrated, and itemized list of the problems impacting on the sysplex's effectiveness in terms of availability, service, workload and performance.

With the degree of complexity involved in managing a sysplexed environment, the exception process can be used as a diagnostic filter to report specific problems. The means of achieving this is by inputting available monitoring sources (that is, the CA MICS files derived from RMF data), comparing this data against predefined exception thresholds, and providing an exception list of potential problems areas.

Exception reporting is designed to reduce the number and volume of reports that the performance analyst has to wade through for detection of potential bottlenecks in the sysplex. The key requirements for successful use of this analysis are setting of appropriate thresholds, and regular, periodic, execution and monitoring of the generated reports.

2.8.3 Standard Output

The sysplex analysis reports give you the capability of tracking the operation of a parallel sysplex. The reports provide data at a number of different levels of the sysplex structure encompassing both the systems, themselves, and the coupling facility (CF). You can use the inquiry options described in Section 2.8.5 to limit reporting to specific systems, times, and members.

This section describes the report formats of the five reports that define the analysis function. The output of these reports is described in the following subsections:

- 1 - Common Report Elements
- 2 - Sysplex Configuration Map
- 3 - Sysplex XCF Member Profile
- 4 - Sysplex Performance Exceptions
- 5 - Sysplex XCF Channel Activity
- 6 - Coupling Facility Configuration

2.8.3.1 Common Report Elements

This section describes the elements common to most of the sysplex analysis reports. Elements that are unique to a specific report are documented with that report.

TITLE ELEMENTS

Each report has up to five lines of title information that give the name of the report, the installation name, the run date, the date(s) being reported, and usually some sort of identification of the Coupling Facility or Member being reported on. The following elements may appear in the standard title area of each report:

Installation Name:	The installation name as defined by the CA MICS NAME parameter within the SITE member that was created at CA MICS installation time. This value is centered in the first line of the report. For more information about this item, consult your CA MICS System Administrator.
SYSTEM:	The CA MICS logical SYSID being reported.
XCF SYSTEM:	The system name as specified by the SYSNAME parameter in the SYS1.PARMLIB member IEASYSxx.
XCF MEMBER:	A specific function, or instance, of a multisystem application. An XCF Member resides on one system and can communicate with other members of the same XCF group across the sysplex.
XCF GROUP:	A unique name for a set of related XCF Members that a multisystem application defines to XCF.
XCF MEMBER RES:	The name of of XCF system on which a specific XCF member resides.
CF SYSTEM:	The system name as specified by the SYSNAME parameter in the SYS1.PARMLIB member IEASYSxx.
COUPLING FACILITY NAME:	The name of the coupling facility being reported.
CF STRUCTURE:	The name of the structure allocated in this coupling facility. Structures are used to implement high-speed serialization and data

sharing.

ZONE: If data is being reported by zone, the zone number is shown.

Run Date: The report execution date in SAS WORDDATE format.

Report Date: The date or date range of the data displayed on this page of the report.

ADDITIONAL ELEMENTS

Reporting Interval: The first row of some reports shows the date or date/time combination of the period being reported. If the data being reported on is from the WEEKS timespan, the date reflects the starting date of the week, as derived from the variable STARTTS. In all other cases, the date reflects the ending date, derived from the variable ENDTS. In monthly reports, only the year and month portions of the date are displayed.

% CHANGE: The last column on some reports, this field is the percent change for each report item (row) between the base period (first data column) and the last period (last data column). If data for only one period is reported, then dashes (---) are displayed. If the base period is zero, then N/A is printed to indicate that a percent change could not be calculated.

2.8.3.2 Sysplex Configuration Map

The Sysplex Configuration Map allows you to address the structure of the sysplex in terms of what members are defined, and what sort of signaling activity between systems is occurring. Because the report is designed to identify intersystem activity, XCF members that confine themselves to a single system are not shown, and members with no signaling activity (no signals sent or received) are also eliminated. Figure 2-210 shows a sample Sysplex Configuration Map.

REPORT OPTIONS

A single report encompassing the entire sysplex is produced. The report can be narrowed to address specific systems, XCF groups, and XCF Member in the sysplex by using the analysis selection options. These and other selection options are discussed in Sections 2.8.5 and 2.8.5.1.

DATA ELEMENTS

A report is produced covering the maximum duration of the input data. The System Configuration Map contains the following information:

INTERVAL START:	The time and date of the first SCPXCM file observation included in the analysis.
INTERVAL END:	The time and date of the last SCPXCM file observation included in the analysis.
DURATION:	The total number of hours, minutes and seconds summarized by the report.
XCF MEMBER RESIDENCE:	The XCF system on which the reported member(s) reside.
XCF GROUP:	The XCF Group to which the reported member belongs. The group usually defines a set of related members used by one or more multisystem applications.
XCF MEMBER:	The specific member's activity being reported on by this line in the analysis. An XCF member represents a specific function, or instance, or a multisystem application.
XCF SYSTEMS SIGNALLED:	The other systems in the sysplex that had signaling activity from or to the reported

member.

TOTAL SIGNALS SENT: The total number of signal sent from the member to one or more members residing on the system identified in the preceding column.

TOTAL SIGNALS RECEIVED: The total number of signals received by this member from members residing on the XCF system that was signaled.

SIGNALS SENT PER SEC: The signals sent per second from this member to members residing on the XCF system signaled.

SIGNALS RECEIVED PER SEC: The signals received per second by this member from member residing on the XCF system signaled.

XCF CNTRS RESET: The number of times during the duration of the analysis that the XCF counters maintained for this member were reset. Any values here represent a loss of data, usually due to a system failure.

INTERVALS MEMBR ACTIVATED: The number of measurment intervals during the duration of the analysis that the member was activated.

INTERVALS MEMBR DE-ACTIVATED: The number of measurment intervals during the duration of the analysis that the member was deactivated.

CA MICS Performance Manager										1
SYSPLEX CONFIGURATION MAP										17:55 Tuesday, January 24, 1995
INTERVAL START: 23JAN95:14:01										
INTERVAL END : 23JAN95:17:00										
DURATION : 2:58:20										
XCF MEMBER RESIDENCE	XCF GROUP	XCF MEMBER	XCF SYSTEMS SIGNED	TOTAL SIGNALS SENT	TOTAL SIGNALS RECEIVED	SIGNALS SENT PER SEC	SIGNALS RECEIVED PER SEC	XCF CNTRS RESET	INTERVALS MEMBR ACTIVATED	INTERVALS MEMBR DE-ACTIVATED
S076	CJKPLEX	MIMCK076	S072	356433	178218	3960.367	1980.200	0	0	0
	IXCL0002	M115	S072	1066	1047	9.963	9.785	0	0	0
	MREPLEX	MRE76	S072	21336	10668	199.402	99.701	0	0	0
	SYSMCS	S076	S072	64	1102	0.598	10.299	0	0	0
	SYSWLM	S076	S072	172	172	1.607	1.607	0	0	0

Figure 2-210. Sample Sysplex Configuration Map

2.8.3.3 Sysplex XCF Member Profile

The Sysplex XCF Member Profile report is designed to track the usage of signaling resources for one or more XCF members over a period of time. The report is presented from the perspective of the system with which the member is communicating. Each profile represents the member's activity on the reported system. A sample report is shown in Figure 2-211.

REPORT OPTIONS

Ordinarily, you will want to use the analysis selection options provided with this analysis to narrow the report to address specific XCF members of interest. In the absence of such selection criteria, a report is generated for each XCF member on each system, and the resulting output can be quite voluminous. The selection options are discussed in Sections 2.8.5 and 2.8.5.2.

DATA ELEMENTS

The elements reported are divided into two categories:

- o Environment
- o Utilization Measures

ENVIRONMENT

The environment elements describe the characteristics surrounding the sysplex environment in which the reported member was operating.

INTVLS W/ ACTIVATE:	The number of measurement intervals during the reporting period in which this XCF member activated: XCMACTCT.
INTVLS W/ DEACTIVATE:	The number of measurement intervals during the reporting period in which this XCF member deactivated: XCMINACT.
INTVLS XCF CNT RESET:	The number of times during the reporting period that the XCF counters maintained for this member were reset. Any values here represent a loss of data, usually due to a system or member failure of some sort: XCMRSCCT.
TOTAL RMF	The total number of RMF recording intervals in

INTERVALS: the reported time frame: INTERVLS

UTILIZATION MEASURES

The measures reported in this section show both overall member activity on the reported system, as well as the signaling activity for the member being profiled. Keep in mind that the "TOTAL" and "Percent of Total" values reported are total for all members communicating with the reported system; not for the sysplex as a whole.

SIGNALS/SEC RECEIVED:	The signals received per second by this member from members residing on the XCF system identified in the "SYSTEM =" portion of the report title: XCMPSRCV
SIGNALS/SEC SENT:	The signals sent per second from this member to members residing on the XCF system identified in the "SYSTEM =" portion of the report title: XCMPSSENT
MEMBER SIGNALS RECD:	The total number of signals received by this member from members residing on the XCF system identified in the "SYSTEM =" portion of the report title: XCMTORCV
MEMBER SIGNALS SENT:	The total number of signals sent from the member to one or more members residing on the XCF system identified in the "SYSTEM =" portion of the report title: XCMTOSNT
TOTAL SIGNALS RECD:	The total number of signals received by members from members residing on the XCF system identified in the "SYSTEM =" portion of the report title. This includes intra-system signaling activity (that is, signals received from other members on the same XCF system).
TOTAL SIGNALS SENT:	The total number of signals sent from all members to members residing on the XCF system identified in the "SYSTEM =" portion of the report title. This includes intrasystem signaling activity (that is, signals sent to other members on the same XCF system).
% OF TOTAL RECD:	The percent contribution of the profiled member to the total signals received.

% OF TOTAL SENT: The percent contribution of the profiled member to the total signals sent.

CA MICS Performance Manager				
HOURLY SYSPLEX XCF MEMBER PROFILE				
SYSTEM = S072 XCF SYSTEM = S072 RUNDATE = Feb 1, yyyy				
XCF GROUP = CJKPLEX XCF MEMBER= MIMCK072 XCF MEMBER RES = S072				

HOUR-WEEK (HR-MM/DD):	14-01/31	15-01/31	16-01/31	CHANGE %

ENVIRONMENT				

INTVLS W/ ACTIVATE	0.00	0.00	0.00	N/A
INTVLS W/ DEACTIVATE	0.00	0.00	0.00	N/A
INTVLS XCF CNT RESET	0.00	0.00	0.00	N/A
TOTAL RMF INTERVALS	1.00	2.00	2.00	100.00 %

UTILIZATION MEASURES				

SIGNALS/SEC RECEIVED	3998.41	3981.19	3946.75	-1.29 %
SIGNALS/SEC SENT	1999.41	1990.64	1973.39	-1.30 %
MEMBER SIGNALS RECVD	67973.00	143323.0	142083.0	109.02 %
MEMBER SIGNALS SENT	33990.00	71663.00	71042.00	109.00 %
TOTAL SIGNALS RECVD	74052.00	228052.0	226188.0	205.44 %
TOTAL SIGNALS SENT	40069.00	228052.0	226188.0	464.49 %
% OF TOTAL RECVD	91.79	62.85	62.82	-31.56 %
% OF TOTAL SENT	84.83	31.42	31.41	-62.97 %

Figure 2-211. Sample Sysplex Member Profile

2.8.3.4 Sysplex Performance Exceptions

The Sysplex Performance Exceptions report is designed to bring to light potential performance bottlenecks and tuning opportunities affecting XCF services and the hardware coupling facility. A sample report is shown in Figure 2-212.

REPORT OPTIONS

The Sysplex Performance Exceptions report has the standard options for selecting specific systems and/or time periods to include. Additionally, the extended options for this report provide for explicitly setting the thresholds that dictate when an exception incident will be reported. The selection options are discussed in Sections 2.8.5 and 2.8.5.3.

REPORT ITEMS

DATE/TIME: The ending time stamp for the period in which an exception was detected, along with the duration of the reporting period. The duration depends on the CA MICS timespan that was used for generating the report.

DESCRIPTION: Text and pertinent instrumentation data concerning the exception, including the measurement item causing the exception, the threshold for the exception, and relevant identification data such as the name of the system or coupling facility involved. The exceptions that can be tracked are discussed in section 2.8.5.3

EXPLANATION/RECOMMENDATION: Text that indicates possible causes for the exception condition, and, where applicable, possible courses of action that might be taken to alleviate or solve the condition.

CA MICS Performance Manager		
SYSPLEX PERFORMANCE EXCEPTIONS		
SYSTEM = S072 RUNDATA = Feb 1, yyyy XCF SYSTEM = S072		
DATE/TIME	DESCRIPTION	EXPLANATION/RECOMMENDATION
31JAN95:16:30:00.00 DURATION: 2:28:20	Requests encountering lock contention exceeds threshold. Requests delayed: 166 threshold: 100 Name of System using CF: S072 Coupling Facility Name: COUPLER CF Structure Name: JESCF	Lock structures are used by CICS and IMS to implement data sharing and high-speed serialization. Application tuning and/or placement may be needed to reduce locking requirements.
31JAN95:17:00:00.00 DURATION: 0:30:00	Requests encountering lock contention exceeds threshold. Requests delayed: 36 threshold: 30 Name of System using CF: S072 Coupling Facility Name: COUPLER CF Structure Name: JESCF	Lock structures are used by CICS and IMS to implement data sharing and high-speed serialization. Application tuning and/or placement may be needed to reduce locking requirements.
31JAN95:16:30:00.00 DURATION: 2:28:20	XCF Messages migrated due to no path exceeds threshold. Messages migrated: 8873 Threshold specified: 1500 XCF Sys Name: S076 Transport Class: TCMIM Message Traffic Direction: OUTBOUND	Most likely, there are no signaling paths assigned to the selected transport class, incurring additional XCF, overhead for selecting alternate transport classes. Review signaling path assignments for the transport class shown.
31JAN95:17:00:00.00 DURATION: 0:30:00	XCF Messages migrated due to no path exceeds threshold. Messages migrated: 1795 Threshold specified: 1500 XCF Sys Name: S076 Transport Class: TCMIM Message Traffic Direction: OUTBOUND	Most likely, there are no signaling paths assigned to the selected transport class, incurring additional XCF, overhead for selecting alternate transport classes. Review signaling path assignments for the transport class shown.

Figure 2-212. Sample Sysplex Performance Exceptions

2.8.3.5 Sysplex XCF Channel Activity

The Sysplex XCF Channel Activity report allows you to track the utilization of the XCF channels linking multiple systems through the coupling facility. A sample report is shown in Figure 2-213.

REPORT OPTIONS

The Sysplex XCF Channel Activity report supports standard time selection options, as well as SYSID and/or XCF System selection. These selection options are discussed in Sections 2.8.5 and 2.8.5.4.

DATA ELEMENTS

The elements reported are divided into two categories:

- o Channel Requests
- o CF Channels

CHANNEL REQUESTS

The Channel Requests elements describe the level of request activity occurring on the reported coupling facility from the perspective of the reporting system (that is, as recorded by RMF running on the reporting system).

TOTAL REQUESTS:	The total number of XCF Channel requests from the reported system to the coupling facility during the time period shown: CFPTOREQ
TOT UNSUCCESSFUL REQS:	The total number of unsuccessful XCF Channel requests from the reported system to the coupling facility during the time period shown: CFPTOURQ
REQUESTS/SECOND:	The average number of requests per second from the reported system to the coupling facility during the time period shown: CFPPSREQ
UNSUCCESS REQS/SECOND:	The average number of unsuccessful requests per second from the reported system to the coupling facility during the time period shown. This is calculated as CFPTOURQ/DURATION.
AVG UNSUC REQS SVC TM:	The average service time, in microseconds, spent during the reported time period servicing

unsuccessful requests.

CF CHANNELS

These elements display path and subchannel availability measures. These items should be relatively static from one time period to the next unless some sort of reconfiguration activity has occurred.

AVG PATHS AVAILABLE: The average number of paths available to the coupling facility during the reported period: CFPAVPTH.

AVG SUBCHANNEL LIMIT: The average number of subchannels that are available for use during the reported period: CFPAVSCL

AVG SUBCHANNELS DEF: The average number of subchannels defined (via the IOGEN) during the reported time period: CFPAVSCG.

AVG SUBCHANS IN USE: The average number of subchannels that are actually in use during the reported period: CFPAVSCU.

CA MICS Performance Manager						
HOURLY SYSPLEX XCF CHANNEL ACTIVITY						
SYSTEM = S072 CF SYSTEM = S072						
COUPLING FACILITY NAME = COUPLER RUNDATE = Feb 1, yyyy						

HOURLY-WEEK (HR-MM/DD):	14-01/28	15-01/28	16-01/28	17-01/28	CHANGE %	

CHANNEL REQUESTS						

TOTAL REQUESTS	116640.0	252185.0	250002.0	124873.0	7.05 %	
TOT UNSUCCESSFUL REQS	0.00	0.00	0.00	0.00	N/A	
REQUESTS/SECOND	68.61	70.05	69.45	69.37	1.10 %	
UNSUCCESS REQS/SECOND	0.00	0.00	0.00	0.00	N/A	
AVG UNSUC REQS SVC TM	0.00	0.00	0.00	0.00	N/A	

CF CHANNELS						

AVG PATHS AVAILABLE	1.00	1.00	1.00	1.00	---	
AVG SUBCHANNEL LIMIT	2.00	2.00	2.00	2.00	---	
AVG SUBCHANNELS DEF	2.00	2.00	2.00	2.00	---	
AVG SUBCHANS IN USE	2.00	2.00	2.00	2.00	---	

Figure 2-213. Sample Sysplex XCF Channel Activity Report

2.8.3.6 Coupling Facility Configuration

The Coupling Facility Configuration report tracks the coupling facility activity in terms of utilization, structure configuration and request activity. A sample report is shown in Figure 2-214.

REPORT OPTIONS

The Coupling Facility Configuration report supports standard time selection options, SYSID and/or XCF System selection, and specific XCF Structure selection. These selection options are discussed in Sections 2.8.5 and 2.8.5.5.

DATA ELEMENTS

The elements reported are divided into three categories:

- o CF CPU Activity
- o CF Configuration
- o CF Request Activity

CF CPU ACTIVITY

The CPU Activity elements describe the utilization of processing power on the coupling facility itself. These measures are repeated for each structure that was selected for reporting.

AVG PROCESSORS:	The average number of processors on this Coupling Facility during the time period shown: CFPAVCPU.
AVG % PROCESSOR BUSY:	The average percent busy for the processors on this Coupling Facility during the time period shown: CFPPCBSY.
AVG % PROCESSOR WAIT:	The average percent wait time for the processors on the Coupling Facility during the time period shown: CFPPCWAT.
PERCENT CF CONNECTED:	The percentage of intervals (RMF recording intervals) during which the Coupling Facility was connected to the reporting system at the end of the measurement interval: CFPPCCON

MAX % PROCESSOR BUSY:	The maximum percent busy for the processors on the Coupling Facility during the time period shown: CFPMXPCB

MAX % PROCESSOR WAIT:	The maximum percent wait time for the processors on the Coupling Facility during the time period shown: CFPMXPCB

CF CONFIGURATION

These elements display storage available and allocated information for the coupling facility and reported structure. These items should be relatively static over the duration of the analysis unless some sort of reconfiguration activity has occurred.

AVG STORAGE DEFINED:	The average amount of storage defined on the coupling facility. This includes both control and non-control storage: CFPVTSDF.

AVG STORAGE FREE:	The average amount of free storage on the Coupling Facility during the time period shown. This includes both control and non-control storage: CFPVTSF.

AVG CNTL STG DEFINED:	The average amount of control storage defined on the Coupling Facility during the time period shown. The difference between this value and the average storage defined represents the average non-control storage: CFPVTSDF.

AVG CNTL STG FREE:	The average amount of free control storage on the Coupling Facility during the time period shown: CFPVCSF.

AVG STRUCTURE SIZE:	The average size of the reported structure during the time period shown: CFSVSI.

CF REQUEST ACTIVITY

The elements reported in this section reflect the request activity associated with the specific structure defined to the coupling facility and being reported on in the analysis.

% STRUCTURE CONNECT:	The percentage of intervals (RMF recording intervals) during the reported time period, in

	which the structure was connected to the system at the end of the interval: CFRPCCON
ASYNC REQS/SEC:	The average number of asynchronous requests per second during the reported time period from the reporting system using the Coupling Facility structure: CFRPSARC.
SYNCH REQS/SEC:	The average number of synchronous requests per second during the reported time period from the reporting system using the Coupling Facility structure: CFRPSSRC.
QUEUED REQS/SEC:	The average number of queued requests per second during the reported time period from the reporting system using the Coupling Facility structure: CFRPSQRC.
REQS/SEC COMPLETED:	The average number of requests completed per second during the reported time period originating on the reporting system: CFRPSRQC.
DUMP DELAYED REQS/SEC:	The average number of request per second that were delayed due to dump serialization activity on the Coupling Facility: CFRPSDRC.
AVG SVC TIME ASYNCH:	The average service time for handling the asynchronous request activity from the reporting system using this Coupling Facility structure: CFRAVATM.
AVG SVC TIME SYNCH:	The average service time for handling the synchronous request activity from the reporting system using this Coupling Facility structure: CFRAVSTM.
AVG SVC TIME QUEUED:	The average service time for handling the queued request activity from the reporting system using this Coupling Facility structure: CFRAVQTM.
DUMP SER DELAY RATE:	The average number of Coupling Facility requests per second, that were delayed because dump serialization was in process during the time period shown: CFRPSDUQ
HI PRI REQ RATE:	The average number of Coupling Facility high priority requests that were queued, per second,

during the time period shown: CFRPSHPQ

LO PRI The average number of Coupling Facility low
REQ RATE: priority requests that were queued, per second,
 during the time period shown: CFRPSLPQ

CA MICS Performance Manager

HOURLY COUPLING FACILITY CONFIGURATION

SYSTEM = S072 CF SYSTEM = S072 RUNDATE = Feb 2, yyyy
COUPLING FACILITY NAME = COUPLER CF STRUCTURE= IXCXCFL

HOURLY DATE (HR-MM/DD):	14-08/yy	15-08/yy	16-08/yy	17-08/yy	CHANGE %

CF CPU ACTIVITY					

AVG PROCESSORS	2.00	2.00	2.00	2.00	---
AVG % PROCESSOR BUSY	28.11 %	28.79 %	28.47 %	28.42 %	1.10 %
AVG % PROCESSOR WAIT	71.87 %	71.19 %	71.52 %	71.56 %	-0.43 %
PERCENT CF CONNECTED	100.0 %	100.0 %	100.0 %	100.0 %	---
MAX % PROCESSOR BUSY	28.11 %	28.84 %	28.48 %	28.42 %	1.10 %
MAX % PROCESSOR WAIT	71.87 %	71.25 %	71.53 %	71.56 %	-0.43 %

CF CONFIGURATION					

AVG STORAGE DEFINED	240M	240M	240M	240M	---
AVG STORAGE FREE	218M	218M	218M	218M	---
AVG CNTL STG DEFINED	240M	240M	240M	240M	---
AVG CNTL STG FREE	218M	218M	218M	218M	---
AVG STRUCTURE SIZE	6308K	6308K	6308K	6308K	---

CF REQUEST ACTIVITY					

% STRUCTURE CONNECT	100.0 %	100.0 %	100.0 %	100.0 %	---
ASYNCH REQS/SEC	62.058	61.683	61.177	61.296	-1.22 %
SYNCH REQS/SEC	0.000	0.000	0.000	0.000	N/A
QUEUED REQS/SEC	0.279	0.220	0.222	0.229	-17.92 %
REQS/SEC COMPLETED	62.058	61.683	61.177	61.296	-1.22 %
DUMP DELAYED REQ/SEC	0.000	0.000	0.000	0.000	N/A
AVG SVC TIME ASYNCH	0.007	0.007	0.007	0.007	---

AVG SVC TIME SYNCH	0.000	0.000	0.000	0.000	N/A	
AVG SVC TIME QUEUED	0.002	0.003	0.003	0.003	50.00 %	
DUMP SER DELAY RATE	0.000	0.000	0.000	0.000	N/A	
HI PRI REQ RATE	0.000	0.500	0.500	0.000	N/A	
LO PRI REQ RATE	0.000	0.000	0.000	0.000	N/A	

Figure 2-214. Sample Coupling Facility Configuration Report

2.8.4 Parallel Sysplex Concepts

This section is an introduction to the concepts underlying the parallel sysplex and its activity.

The following sections give an overview of parallel sysplex concepts and the controls available:

- 1 - Sysplex Functional Overview
- 2 - Controlling the Sysplex Environment

2.8.4.1 Sysplex Functional Overview

IBM's System/390 Parallel Sysplex offering represents a significantly different approach to data processing, and an exciting step in the continuing evolution of computing. Hardware and software combine to provide support for full data sharing among the processors in the sysplex and dynamic workload balancing across the systems.

With the introduction of Sysplex, there also comes a whole array of terminology necessary for quantifying the capabilities and components of the technology. Below is a brief definition of some of the major terms:

- o **SYSplex (SYStem comPLEX):** As the name would indicate, this represents a strategy for linking together (coupling) multiple MVS images with a standard communication interface. The idea is to ultimately provide a single point of control for all processors in the organization. At a somewhat lower level, a sysplex can be defined more concretely as a set of MVS systems that have been initialized with the same sysplex name and are fully connected by XCF (see below) signaling paths.
- o **Coupling Facility:** Hardware connecting two or more System/390 systems, which allows multiple MVS/ESA system images to do workload balancing. The Coupling Facility provides support for serialization and buffer management, caching for shared data, as well as message queuing to ensure the integrity and consistency of data throughout the Parallel Sysplex environment. Communication between coupled systems is via very-high-speed fiber optic links that can extend up to three kilometers (almost 2 miles).
- o **SYSplex Timer:** An external hardware timer shared by all of the systems in a sysplex, which effectively moves the time of day clock control from the individual processors to a common external clock capability. This solves a long standing issue in the multisystem environment of dealing with cycle synchronization, creating a common time stamp across the multiple systems composing the sysplex.
- o **XCF (Cross System Coupling Facility):** Not to be confused with the Coupling Facility, which is hardware as described above, XCF represents multisystem services provided as part of the operating system to communicate to authorized programs on other systems. The task of managing multiple systems connected in a sysplex is similar to the challenge of multitasking in a single system. The major additional

requirement is a vehicle for communications between the systems, and XCF is that vehicle.

- o Sysplex Couple data set: MVS requires a DASD data set to be shared by all systems in the sysplex. In this data set, MVS stores information related to the sysplex systems, XCF groups, and their members.
- o Multisystem Application: An application program that has distributed functions across MVS systems. GRS (Global Resource Serialization) is an example of such an application program. A multisystem application can be an installation-defined program, an MVS component or subsystem, or a program product.
- o XCF MEMBER: A specific function, or instance, of a multisystem application. An XCF Member resides on one system and can communicate with other members of the same group (see below) across the sysplex.
- o XCF GROUP: A unique name for a set of related XCF Members that a multisystem application defines to XCF. Communication between group members on different systems occurs over the signaling paths that connect the systems; on the same system, communication between group members occurs through local signaling services. As an example, SYSGRS is the XCF group name for the the global resource serialization application.
- o Transport Class: A means of associating one or more XCF groups (based on similar signaling requirements) and then assigning them signaling resources (signaling paths and message buffers). A transport class allows you to segregate message traffic according to the needs of an applications's group, the lengths of its messages, or both.

The goal of the Parallel Sysplex is to provide for horizontal growth of a computing environment while beginning the task of simplifying the operation of that environment. Prior to sysplex, each system application had to handle issues of communication with other systems in a multisystem environment. The introduction of XCF removes from system applications the redundant communication functions, allowing those applications to utilize XCF in a high level manner.

So, instead of depending on a communications vehicle like VTAM for cross-system services, XCF has a direct, fully configured, connection to all systems in the complex. The

major users of XCF services are the IBM system functions that need to communicate with other systems. Examples include GRS (Global Resource Serialization) and console support. Additionally, in support of workload balancing, major transaction subsystems such as CICS and IMS utilize XCF services to shift transaction processing dynamically to the appropriate region, regardless of the system from which the transaction originated.

2.8.4.2 Controlling the Sysplex Environment

While there are a number of SYS1.PARMLIB members that define and affect a sysplex, the primary means of controlling a sysplex is through the specifications made in the COUPLExx PARMLIB members, and through the establishment of "policies" that control management decisions spanning two or more systems.

The COUPLExx members contain the critical definitions of Signaling Paths and Transport Classes. XCF Groups (groups of related members) should be assigned to transport classes based on message lengths, frequency, and performance requirements.

The policies that control a sysplex are not defined in SYS1.PARMLIB, but instead, each of the needed policies reside in a couple data set.

A policy is a set of rules and actions that govern the behavior of systems in a sysplex when using certain MVS services. Each policy provides a means for MVS to manage specific resources in compliance with your workload and system requirements. A policy can be set up for just a single system, several systems, or all systems in the sysplex.

There are three types of policies that can be defined. You can define more than one policy of each type, but only one policy of each type can be active at a time. For example, you might define a workload management policy (one of the three types) for prime shift, and another for off-prime shift.

The following three policy types can be used to control the system management in a sysplex:

- o The coupling facility resource management (CFRM) policy defines rules and actions for MVS management of coupling facility resources.
- o The sysplex failure management (SFM) policy determines how MVS is to manage system failures, signaling connectivity failures, and actions that reconfigure PR/SM.
- o The workload management (WLM) policy allows you to define service goals for your various workloads (for example, batch, TSO, CICS.)

Workload management policies are not presently addressed by the CA MICS Performance Manager Option Sysplex Analysis. Consequently, the discussion here is limited to the following topics:

- 1 - Planning for Signaling Services
- 2 - Coupling Facility Resource Management
- 3 - Sysplex Failure Management

2.8.4.2.1 Planning for Signaling Services

Signaling is the means by which XCF group members communicate in a sysplex environment. Signaling can be achieved through use of Channel-to-Channel (CTC) connections, a coupling facility (using coupling facility list structures), or a combination of the two.

All systems in a sysplex must have full signaling connectivity. That is, an inbound and outbound signaling path must exist between each pair of systems in the sysplex.

Signaling paths defined via CTC connections must be exclusively defined as either inbound or outbound. In contrast, a significant advantage of using coupling facility list structures is that they can be used for both inbound and outbound paths. An added advantage is that the same COUPLExx PARMLIB member can be used for each system in the sysplex.

Transport classes are MVS's way of associating one or more XCF groups (based on similar signaling requirements) and then assigning them resources (paths and message buffers). The idea behind transport class definitions is to segregate message traffic according to XCF group requirements, the lengths of the messages, or both.

As with the paths themselves, transport classes are defined in the COUPLExx member of PARMLIB. The CLASSDEF statements assign a class name, a default buffer length, and XCF groups that will be assigned to the class.

In general, you want to specify a class length (CLASSLEN keyword) equal to the length of the messages most frequently sent by the group(s) in the transport class. Selecting an appropriate class length represents the classic tradeoff between storage used and performance for signaling activity. If messages are usually shorter than the message buffer, storage is wasted; if longer, performance is degraded.

2.8.4.2.2 Coupling Facility Resource Management

The coupling facility is at the heart of a parallel sysplex environment, and makes possible the sharing of data by providing high-speed locking, caching and message/list services between selected System/390 processors running MVS/ESA SP 5.1 or above.

You manage a coupling facility through a policy. Coupling Facility Resource Management (CFRM) allows you to specify how the coupling facility and attendant resources are to be used. CFRM policy is created using an IBM utility which places the information out on one or more couple data sets. MVS and multisystem applications use the definitions in the CFRM policy when managing the coupling facility storage.

Storage in a coupling facility is divided into different constructs call "structures." These structures are used to implement high-speed serialization and data sharing. To facilitate this process, the structures are divided into three types: cache, list, and lock. In addition to structures, some storage in the coupling facility can be allocated as a dedicated dump space for capturing diagnostic information related to the structures.

Among other things, when defining a CFRM policy, you must identify each coupling facility and structure, and define the amount of space within the coupling facility that you want for each structure. More importantly, for performance and capacity purposes, you can define preference and exclusion lists for each structure in the CFRM. These lists can be defined as follows:

- o A preference list is an ordered list of coupling facility names in which a particular structure can be allocated.
- o An exclusion list is an unordered list of structure names with which the structure being defined should not share a coupling facility.

The system selects a coupling facility for structure allocation based on the preference and exclusion lists in the CFRM policy and the availability of storage in the coupling facility. However, the system will ignore an exclusion list if it would prevent the structure from being allocated.

Keep in mind that for performance and recovery purposes, you can distribute structures on different coupling facilities.

Having more than one coupling facility is almost essential in a sysplex to avoid a single point of failure that would prevent intersystem communication.

2.8.4.2.3 Sysplex Failure Management

SFM policy allows you to specify what actions MVS should take when certain failures occur in the sysplex. Basically this amounts to accommodating the removal of one or more systems from the sysplex in order that the remaining systems can continue to do work.

Two main types of failures come into play in a sysplex environment:

- o System failures, indicated by a status update missing condition
- o Signaling connectivity failures in the sysplex

The first of these is fairly straight forward. If a system does not update its status information within a predetermined period of time, a "status update missing" condition is raised. SFM allows you to specify how to respond to this condition. The most common options for addressing this are to prompt the operator for help in resolving the problem, or to "isolate" the failing system without operator intervention. System isolation terminates I/O and coupling facility accesses, resets channel paths, and loads a nonrestartable wait state on the failing system, thus ensuring that the system is unable to corrupt shared resources.

In the case of a loss of signaling connectivity, the performance issues and policy decisions are more pertinent. All systems in the sysplex must have signaling paths to and from every other system sharing in the sysplex. If, for some reason, signaling connectivity between any sysplex systems is lost, then one or more systems can be removed (under SFM) so that the systems that remain in the sysplex still have full connectivity with each other.

In handling a system connectivity failure, SFM attempts to maximize the aggregate value of the surviving sysplex to the installation. SFM allows you to assign a relative importance to each system in the sysplex. Then, if connectivity between any two systems is lost, SFM decides which of the two systems to retain in the sysplex based upon its relative importance.

How you establish your SFM policy can have a very substantial impact on the performance of the sysplex in the event that some sort of failure does occur. For example, suppose that you have a three system sysplex, SYSA, SYSB, and SYSC.

Further suppose that SYSA is a larger mainframe, and that SYSB is a rather small system, but supporting mission critical applications. If SYSA and SYSB lose connectivity with each other, SFM decides which system to keep in the sysplex, and which one to let go. If SYSA is released, considerable computing power is no longer available to the sysplex, and performance may suffer substantially. If SYSB is released, its critical applications are no longer available, but the sysplex may have enough computing power to continue to provide adequate service in other areas.

Tracking and monitoring sysplex handling of system failures should be an ongoing process, providing feedback to the policymakers who need to decide upon the optimal system failure management configuration.

2.8.5 Component Operation

This Component Operation Section describes the operation of the sysplex analysis routines in terms of the CA MICS MVS Performance Manager inquiry options selected.

Inquiry options available with the sysplex analysis software describe the reports to be generated. In particular, these options define:

- o The cycles from the CA MICS RMF files to be processed.
- o Workload selection by system (SYSID), XCF system, XCF member, XCF group, Coupling Facility structure, date/time, and zone.
- o The summarization level of the reports as hourly, daily, weekly, or monthly. The availability of each summarization level assumes that your system supports the required files at each level.
- o The reports to be produced.

The following fields are common to most control screens:

- Timespan: The timespan (DETAIL, DAYS, WEEKS, or MONTHS) in the CA MICS file(s) you want to use as input for processing.
- Cycle(s): One or more cycles of the CA MICS files that you want to use for analysis. Because SMF data sets are not normally dumped at exactly midnight at most installations, a CYCLE from the CA MICS file will usually contain a 24-hour period spanning two days. Hence, some other type of selection may be required to generate a daily report. This selection capability is provided by the DATE and HOUR control options.
- Database: The single-character CA MICS identifier for the database to be accessed. Valid values are the database unit identifiers available at your installation. Contact your CA MICS System Administrator if you do not know what database units are available, or enter a question mark (?) to view a list of your database units. You can select up to five databases per inquiry.

Allow Run-Time Override:	The indication of whether the database identifier and/or the Cycles specification can be overridden at execution time.
Report by hour:	The indication of whether a report or plot should be generated for each hour.
Generate a report for each value of Zone:	The indication of whether a separate report or plot should be generated for each zone defined in the input data or whether all zones should be merged together for the report or plot.
Date:	Up to two date ranges (DDMONYY format) to be used for analysis. This option is normally used to drop the extra day from a cycle that is due to the timing of the SMF dump.
Hour:	One to three ranges of hours for which data is to be reported. For an individual hour, leave the end hour blank or make it the same as the start hour. You can specify a value from 00 to 23.
Zone:	The specification that zone selection is to be applied. If zone 1 definition is 8:00 a.m. to 6:00 p.m., you can specify reports covering that time period by entering a 1 in the ZONE field.
SYSID:	The specification that SYSIDs are to be selected for reporting. You can enter up to six SYSIDs to be selected. If you leave the field blank, no testing of SYSIDs will be done.
Systems:	The XCF systems to be selected for reporting. If you code this field, only those systems specified will be selected.
Groups:	The XCF member group(s) to be selected for reporting. If you code this field, only those specified will be selected.
Members:	The XCF member(s) to be selected for reporting. If you code this field, only those specified will be selected.
Structures:	The Coupling Facility Structures to be reported. If code this field, only those specified will be selected.

Specify extended options:	The specification that allows you to supply additional report specifications. These options include special SAS data selection statements, nonstandard input and output files, and special graphics options.

The operation of the time selection control parameters depends on definitions made during the installation of CA MICS. If you are not sure about the ZONE definitions for your installation or the CA MICS database update procedures, consult your CA MICS System Administrator.

2.8.5.1 Sysplex Configuration Map

The Sysplex Configuration Map details the members defined in the sysplex, as well as the intersystem signaling activity taking place across the sysplex. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the Sysplex Configuration Map obtains its data from the SCPXCM Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The Sysplex Configuration Map is produced using a series of five screens. Only the first screen is required. The others are used to specify additional options that are normally defaulted.

```

/----- Sysplex Configuration Map -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: SPA001 - Sysplex Analysis - Sample Inquiry |
| Time-span ==> DAYS (DETAIL/DAYS/WEEKS/MONTHS)      Cycle(s) ==> 01 - 02 |
| Data base ==> P (PRIMARY) _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
| Data Selection:                                     Allow run-time override |
| Date          ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Hour          ==> _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ |
| Zone          ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ |
| Sysid         ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ |
| Systems       ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ |
| Groups        ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ |
| Members       ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ |
|                                     _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ |
| Specify extended options ==> N (Y/N) |
|                                     |
\-----/

```

Figure 2-216. Sysplex Configuration Map Inquiry Screen

Figure 2-216 shows a Sysplex Configuration Map inquiry screen. The fields common to all Sysplex Analysis routines are described in Section 2.8.5.

There are four other screens in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.8.5.2 Sysplex XCF Member Profile

The Sysplex XCF Member Profile Report tracks the signaling resource usage for one or more XCF members over a period of time. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database.

Note: The Sysplex XCF Member Profile report obtains its data from the SCPXCM Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The Sysplex XCF Member Profile report is produced using a series of five screens. Only the first screen is required. The others are used to specify additional options that are normally defaulted.

```

/----- XCF Member Profile -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: SPA001 - Sysplex Analysis - Sample Inquiry |
| Time-span ==> DAYS (DETAIL/DAYS/WEEKS/MONTHS)   Cycle(s) ==> 01 - 02 |
| Data base ==> A (PRIMARY) _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
| Report by hour if timespan is DAYS or WEEKS ==> Y (Y/N) |
| Generate a report for each value of Zone ==> N (Y/N) |
| Data Selection:                                     Allow run-time override |
| Date ==> _ _ _ _ _ - _ _ _ _ _ - _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Hour ==> _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ - _ _ _ _ _ _ _ |
| Zone ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Sysid ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Systems ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Groups ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Members ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-217. Sysplex XCF Member Profile Inquiry Screen

Figure 2-217 shows a Sysplex XCF Member Profile inquiry screen. The fields common to all sysplex analysis routines are described in Section 2.8.5. Generally speaking, you will want to use the Member or Group selection criteria to narrow the focus of your analysis to selected members or member groups of interest. The Sysplex Configuration Map described in the previous section can be used as a starting point in determining which members and/or groups must be profiled using this facility. The following are general comments on this screen.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Report by hour and Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized to the timespan level (that is, MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

There are four other screens in this series that you can use to specify additional options that are normally defaulted. They are the following:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus, see Section 1.4.3.6, Specifying Common Extended Options.

2.8.5.3 Sysplex Performance Exceptions

The Sysplex Performance Exceptions report allows you to detect and track various unusual or out-of-the-ordinary conditions that might arise in XCF service processing and the underlying coupling facility hardware configuration. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the Sysplex Performance Exceptions report obtains its data from the SCPXYC, SCPXCA, HARCFF, and HARCFR Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The Sysplex Performance Exceptions report is produced using a series of six screens. Only the first screen is required. The others are used to specify additional options that are often defaulted.

```

----- XCF Performance Exception Analysis -----
| Command ==>
|
| Composing CA MICS Inquiry: SPA001 - Sysplex Analysis - Sample Inquiry
|
| Time-span ==> DAYS    (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 02
| Data base ==> A (PRIMARY) _ _ _ _ _                      _ _ _ _ _
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N)
|
| Data Selection:                                           Allow run-time override
|   Date           ==> _ _ _ _ _ - _ _ _ _ _      _ _ _ _ _ ==> N (Y/N)
|   Hour           ==> _ _ - _ _      _ _ - _ _      _ _ - _ _ ==> N (Y/N)
|   Zone           ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> N (Y/N)
|   Sysid          ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> N (Y/N)
|   Systems        ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _
|
| Specify extended options ==> N (Y/N)

```

Figure 2-218. Sysplex Performance Exceptions Inquiry Screen

Figure 2-218 shows a Sysplex Performance Exceptions inquiry screen. The fields common to all Sysplex Analysis routines

are described in Section 2.8.5.

There are five other screens in this series that you can use to specify additional options that are often defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Exception Thresholds
- o Alternate Input File
- o Save Selected Input Observations

The Exception Thresholds Extended Option is unique to this query. A description of how to set the Exception Thresholds follows. For additional information on the common Extended Options refer to Section 1.4.3.6, Specifying Common Extended Options.

If you enter Y (Yes) in the Specify extended options field, the Extended Options Menu screen, shown in Figure 2-219, is displayed.

```
/----- Extended Options Menu -----\
| Command ==>
|
| Inquiry Step: XCF Performance Exceptions
|
| Specify global data selection exit          ==> N (Y/N/R)
| Specify Exception Thresholds                ==> N (Y/N/R)
|
| Specify alternate input file                 ==> N (Y/N/R)
| Save selected input observations             ==> N (Y/N/R)
|
|
|
|
|
|
|
|
|
|
```

Figure 2-219. Sysplex Performance Exceptions Extended Options Menu Screen

If you enter Y (Yes) to the Specify Exceptions Thresholds in the Extended Options Menu, the Exception Thresholds screen, shown in Figure 2-220 is displayed. This screen allows you to enter trigger values for generating exception incidents when the analysis is run.

```

----- Exception Thresholds -----
| Command ==>
|
| Inquiry step: XCF Performance Exceptions
|
|-----Threshold Values-----
| Exceptions:          DETAIL      DAYS      WEEKS      MONTHS
| Intervals w/ XCF Deactivation    0          1          3          5
| XCF Count Resets                 0          1          3          5
| Pct XCF Messages Larger than Buffer 5          5          5          5
| XCF Messages Rejected            0          0          0          0
| XCF Messages Migrated - No Path  0          1          3          5
|
| Pct XCF Path Busy                50         50         50         50
| Intervals w/ XCF Path Deactivation 0          1          3          5
| Inbound Signals Refused           0          0          0          0
|
| CF Pct Processor Busy            5          5          5          5
| Unsuccessful CF Requests          0          1          3          5
| CF Requests failed - Path Busy    0          1          3          5
|
| CF Requests delayed due to dump ser 0          1          3          5
| CF Requests hitting lock contention 0          1          3          5
|
|
|
|-----

```

Figure 2-220. Sysplex Performance Exceptions Threshold Setting Screen

When using the XCF Performance Exception Analysis, it is recommended that you start out with the default exceptions shown above, and then adjust them to reflect the characteristics of your particular environment. The four columns correspond to the four CA MICS timespans that can provide input to the analysis. Only the column that corresponds to the timespan selected on the primary screen will be used in the analysis. For example, if the timespan selected on the primary screen was WEEKS, the threshold values in the WEEKS column will be the ones used in the analysis.

The length of RMF intervals, and the definition of CA MICS ZONES play an important role in setting appropriate exception threshold values. At the DETAIL level, a 5 or 10 minute RMF interval might dictate a set of exception values considerably different than a 20 or 30 minute interval. Similarly, at the WEEKS and MONTHS levels, each CA MICS observation contains ZONE in the file key, and this determines the overall duration. For example, at the MONTHS timespan, ZONE 1 might cover prime shift (9 to 5 Monday through Friday) for an entire month. If ZONE 2 covered the remaining 16 hours of each day Monday through Friday, the observations describing ZONE 2 would cover a period of time that is twice as long as those for ZONE 1. Because of this, you should probably use ZONE selection criteria and separate analyses for each ZONE when dealing with WEEKS or MONTHS level data.

Keep in mind, that if you set up an analysis using one timespan as input along with corresponding threshold values established via this screen, and then later change the analysis to utilize a different timespan, the corresponding column of threshold values specified here will go into effect. Changing timespans on the primary screen does NOT reset the exception threshold values to the defaults.

The grouping of exceptions shown in this screen is indicative of the CA MICS file from which the exception is generated. Thus, the first five exceptions are from the SCPXCY file, the next three are from the SCPXCA file, the next three from the HARCFCP file, and the last two from the HARCFCR file.

To turn off any give exception, simply set the threshold value to an artificially high value.

2.8.5.4 Sysplex XCF Channel Activity

The Sysplex XCF Channel Activity report allows you to track the utilization of the XCF channels linking multiple systems through the coupling facility. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS timespans if they are present in your CA MICS database. Note that the Sysplex XCF Channel Activity report obtains its data from the HARCFP Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The Sysplex XCF Channel Activity report is produced using a series of five screens. Only the first screen is required. The others are used to specify additional options that are normally defaulted.

```

----- XCF Channel Activity -----
| Command ==>
|
| Composing CA MICS Inquiry: SPA001 - Sysplex Analysis - Sample Inquiry
|
| Time-span ==> DAYS    (DETAIL/DAYS/WEEKS/MONTHS)          Cycle(s) ==> 01 - 02
| Data base ==> A (PRIMARY) _ _ _ _ _                      _ _ _ _ _
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N)
|
| Report by hour if timespan is DAYS or WEEKS ==> Y (Y/N)
| Generate a report for each value of Zone      ==> N (Y/N)
|
| Data Selection:                                           Allow run-time override
|   Date           ==> _ _ _ _ _ - _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> N (Y/N)
|   Hour           ==> _ _ - _ _ _ _ - _ _ _ _ - _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> N (Y/N)
|   Zone           ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> N (Y/N)
|   Sysid          ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ ==> N (Y/N)
|   Systems        ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _
|
| Specify extended options ==> N (Y/N)
|
|
|

```

Figure 2-221. Sysplex XCF Channel Activity Inquiry Screen

Figure 2-221 shows a Sysplex XCF Channel Activity inquiry screen. The fields common to all Sysplex Analysis routines are described in Section 2.8.5. The following are general comments on this screen.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Report by hour and Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized to the timespan level (that is, MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

There are four other screens in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.8.5.5 Coupling Facility Configuration

The Coupling Facility Configuration Report tracks the coupling facility activity in terms of utilization, structure configuration and request activity by structure name. It can be generated for the DETAIL, DAYS, WEEKS, or MONTHS time-spans if they are present in your CA MICS database. Note the Sysplex XCF Structure & Activity report obtains its data from the HARCFS, HARCPR and HARCPR Files. If you are unsure of the timespans supported at your installation, consult your CA MICS System Administrator.

The Coupling Facility Configuration report is produced using a series of five screens. Only the first screen is required. The others are used to specify additional options that are normally defaulted.

```

/----- Coupling Facility Configuration -----\
| Command ==>                                     |
| Composing CA MICS Inquiry: SPA001 - Sysplex Analysis - Sample Inquiry |
| Time-span ==> DAYS (DETAIL/DAYS/WEEKS/MONTHS)      Cycle(s) ==> 01 - 02 |
| Data base ==> A (PRIMARY) _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Allow Run-Time Override of Data Base ==> N (Y/N) and/or Cycle(s) ==> N (Y/N) |
| Report by hour if timespan is DAYS or WEEKS ==> Y (Y/N) |
| Generate a report for each value of Zone ==> N (Y/N) |
| Data Selection:                                     Allow run-time override |
| Date ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Hour ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Zone ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Sysid ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Systems ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Structures ==> _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 2-222. Coupling Facility Configuration Inquiry Screen

Figure 2-222 shows a Coupling Facility Configuration inquiry

screen. The fields common to all sysplex analysis routines are described in Section 2.8.5. Generally speaking, you'll want to use the Structure selection criteria to narrow the focus of your analysis to selected structures. The following are general comments on this screen.

In addition to the timespan qualifier, you have the option of selecting summaries by ZONE and HOUR. Summarization by HOUR can be done if you are using the DAYS or WEEKS timespans. Summarization can be specified with the Report by hour and Generate a report for each value of Zone options.

If you specify the DAYS timespan and summarization by HOUR, and if you request multiple cycles, a separate report is produced for each day. If you do not specify a summarization parameter, the report is summarized to the timespan level (that is, MONTHS is summarized to the month, WEEKS to the week, DAYS to the day, and DETAIL to the RMF interval).

There are four other screens in this series that you can use to specify additional options that are normally defaulted. They are:

- o Extended Options Menu
- o Global Data Selection
- o Alternate Input File
- o Save Selected Input Observations

For a description of the Extended Options Menus refer to Section 1.4.3.6, Specifying Common Extended Options.

2.9 WLM Goal Mode Analysis

This section introduces a series of queries that analyze WLM Goal Mode workload measurements to help determine if service objectives are being met and to provide a first-level analysis of throughput problems.

Workloads can be summarized by WLM Service Class, Importance, Workload, or Report Class. Measurements involving both resource utilization and throughput are included. Additionally, a separate analysis is provided that combines LPAR processor capacity and resource metrics. Together with WLM Workload measurements, this analysis determines which workloads are having the greatest impact on processor usage and whether processor resources must be added.

These queries are ideal in determining which workload and time period would be most beneficial to further analyze with modeling solutions such as the z/OS performance modeling solutions delivered by our CA Smart partner; OPNET.

The following sections describe the operation of the WLM Goal Mode Service Analysis inquiries and how to use the reports and graphs produced:

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output

2.9.1 Functional Description

The WLM Goal Mode Service Analysis inquiries allow you to track and analyze the WLM workload resources and throughput measurements, and to perform first-level problem resolution analysis against performance problems. Any supported CA MICS timespan and summarization pattern can be implemented, a series of tabular reports can be produced, and a CSV file that is compatible with CA MICS Q&R can be produced for subsequent generation of charts and graphs.

The WLM Service Definition report provides a report of the complete service definitions used to generate the WLM usage data selected. This report is useful for better understanding the output produced from the other WLM Service Analysis queries.

This query uses the following CA MICS files as input:

WLMSEA - Service Class
WLMSEDE - Service Definition

The WLM Service Delivery reports provide a summary of WLM services delivered, WLM service objectives, and resources consumed, either by SYSPLEX or LPAR. Additionally, the summarization may be performed by WLM Service Class, Importance level, Workload, or Report Class.

These queries use the following CA MICS files as input:

WLMSEC - Service Class Resource Consumption or
WLM_EC - Report Class Resource Consumption

WLMSED - Service Class General Execution Delay or
WLM_ED - Report Class General Execution Delay

WLMSEA - Service Class
WLMSEDE - Service Definition

WLMSEC - Service Class Resource Consumption
WLM_EC - Report Class Resource Consumption

The LPAR Service Delivery reports provide a summary of LPAR processor resource usage, WLM service delivered, and workload resources consumed, by LPAR. Additionally, the summarization may be performed by WLM Service Class, Importance level, Workload, or Report Class.

These queries use the following CA MICS files as input:

HARCPU - CPU Activity File
WLMSEC - Service Class Resource Consumption or
WLM_EC - Report Class Resource Consumption

WLMSED - Service Class General Execution Delay or
WLM_ED - Report Class General Execution Delay

WLMSEA - Service Class
WLMSEDE - Service Definition

WLMSEC - Service Class Resource Consumption
WLM_EC - Report Class Resource Consumption

The range of data that can be reported from a particular CA MICS database depends on the timespans that were defined for these files and the number of cycles defined in each timespan. If you are not familiar with the generation options for the database unit containing the MVS HAR or WLM Information Areas, ask your CA MICS System Administrator for this information.

2.9.2 Usage Guidelines

The WLM Goal Mode Service inquiries are a series of reports that provide an overview of how service is being delivered and identify potential delays to service delivery. These reports let you quickly determine if service objectives are being met during the reporting intervals chosen and perform a first-level problem analysis when objectives are not being met.

The reports are grouped according to usage:

- o WLM Service Definition Entries

This is a detailed listing of the WLM Service Definitions that were in effect during the reporting interval. A good understanding of the service definitions is required to properly interpret the other reports. They can help explain why a particular service objective is not being met.

- o Analysis of Service Delivered, by WLM Classes

These reports provide a summary of delivered service by SYSPLEX, and optionally by LPAR. The reports contain data on SRM-based transaction service, resources consumed, and delays. They can be ordered by WLM Service Class and Period, WLM Report Class, WLM Importance, or WLM Workload group.

- o Analysis of Service Delivered by LPAR

These reports provide a summary of service by LPAR, including processor usage in times and MSUs, together with all of the workloads contributing to the LPAR busy. This helps you understand and quantify those workloads that are driving the LPAR consumption, so you can plan for use of the Defined Capacity Limit function required to implement the sub-capacity licensing feature. The reports are ordered by LPAR. The workload information can be summarized by WLM Service Class, WLM Report Class, WLM Importance, or WLM Workload group.

All of the reports provide support for special purpose processors. In addition to tabular reports, CSVs can be optionally produced for subsequent use by Q&R to create charts, graphs, and reports on the users' workstation.

The following sections discuss new performance metrics that

were added to the CA MICS Performance Manager in support of the WLM Goal Mode Service queries.

- 1 - LP Related Metrics
- 2 - Workload Related Metrics

2.9.2.1 LP Related Metrics

The following metrics were added to enhance the understanding and management of z/OS based LP resources:

Weighted LP Cap. (MSUs)	This is the effect capacity of the LP under study, expressed in MSUs per hour. Because of LP weight enforcement, the actual capacity of an LP to do work is reduced in proportion to its defined weight. Quantifying this value helps in LP sizing and in understanding why performance objectives might not be met.
Weighted LP Share Pct.	The share of active logical processors that is available to this LP, expressed as a percentage. This value is used in computing other metrics within the reports.
No. of Available Procs.	The total number of processors available for use by this LP. This is not a computed value, but rather is the CA MICS data element CPUNOCPU and includes all types of processors configured in the logical partition (standard CP and special processors).
CPs Required for LP Wgt.	The minimum number (or fraction) of CPs needed to ensure that the LP can actually sustain its weighted share. Round this value up to the next integer when defining the LP size. If this value is greater than the number of CPs defined for the LP, the LP will not have sufficient capacity to support its weighted share, and performance may suffer. If this value is significantly lower than the number of CPs defined for the LP, then performance of large processing tasks, such as online application, may be affected.
LP Shr. Proc. Use Ratio	This value is sometimes referred to as "Short CPs." It is a ratio of the MVS Busy time (CPUTOVBT) to the LP Busy time (CPUTOBTM) for a reporting interval. The purpose of this metric is to demonstrate the lack of processor resources in an easy way that can be presented to non-technical personnel. In general, a value greater than 1.25 is cause for concern, especially if performance objectives are not being met.

Weighted CP Speed (MSUs)	The effective speed of an individual CP within the LP configuration. This value is a function of the CPs Required for LP Wgt. metric defined above and the number of CPs active for the LP. The total weight for the LP is spread equally among all active CPs, causing each CP to be dispatched by PRSM for an interval based on the Weighted LP Share percentage mentioned above. Having more CPs active than are required for a given weighted LP capacity will result in the effective CP speed being reduced and can result in performance degradation. IBM's Intelligent Resource Director facility can attempt to control the number of active CPs based on WLM Goal Mode service objectives to help ensure that the number of active CPs is optimal for a given workload.
Percent Ready Tasks Wait	The number of address spaces that are waiting for dispatch on a CP, expressed as a percentage. If the LP Shr Proc. Use Ratio value is high and this value is more than 50%, then a significant processor resource shortage exists and performance may be affected.

2.9.2.2 Workload Related Metrics

CPU Usage Delay Ratio The ratio of CPU Using Samples to the sum of CPU Using and CPU Delay Samples. This is a relative measure of how much the defined workloads are waiting for processor dispatch. When this value is more than 0.50, then the LP Related metrics should be examined closely to better understand how much the performance is being affected by the lack of processor resources.

zAAP Usage Delay Ratio The ratio of IFA Using Samples to the sum of IFA Using and IFA Delay Samples. This is a relative measure of how much the defined workloads are waiting for IFA processor dispatch. Because the zAAP technology is relatively new, no firm guidelines have yet been established for determining when a performance impact might occur. However, given the nature and use of zAAP engines, a value of 0.20 might be cause for concern and could lead you to consider other performance measurements.

While the performance metrics for zAAP usage are basic in nature, understand that the LP weight is used to manage the zAAP engines as well as CPs. Specifying a low LP weight value will not only cause the CPs defined to be minimally dispatched, but will likewise cause the zAAP engines to be minimally dispatched as well. Also, because the zAAP engines are dispatched using the z/OS dispatcher, which must execute on a CP, poor processor throughput will cause Java work to not be dispatched on the zAAP engines, even though sufficient zAAP resources are available.

Additionally, the weight percentages are based on the total weight of the ICF pool and therefore might require a change in the weights for LPs that use ICF engines.

zIIP Usage Delay Ratio The ratio of Samples Using zIIP to the sum of Samples Using zIIP and Samples zIIP Delay. This is a relative measure of how much the defined workloads are waiting for zIIP processor dispatch. Because the zIIP technology is

relatively new, no firm guidelines have yet been established for determining when a performance impact might occur.

2.9.3 Standard Output

COMMON ELEMENTS

This section describes elements common to all WLM Goal Mode Service analysis reports. Options that are unique to a specific report are documented with that report.

TITLE ELEMENTS

Each report has five lines of title information available. All reports give the name of the report, the name of the component, and the run date. In addition, the heading on the System Storage Usage and Private Storage Usage reports details the data selected, the summary level of the report, and the time periods reported by column.

- o For the monthly report, the column title is year-month.
- o For the weekly report, the week is identified by month/day/year. The day is determined from the CA MICS file elements used to generate the report.
- o For the daily report, the day is given as month/day/year.
- o For the hourly level reports, the hour is given in the column title along with the month and day as hour-month/day. The day is determined from the CA MICS file elements used to generate the report.

The following elements are in the standard title portion of each report:

Installation Name:	On some reports, the installation name that is defined at CA MICS system installation time. This value is centered in the first line of the report. For more information about this value, ask your system administrator.
Component Name:	The name of the CA MICS performance component, centered in the first or second line of each page.
Report Name:	The name of the report, centered in the third line of the heading. It cannot be modified.
Sysplex and SYSID:	The sysplex and CA MICS logical SYSID being reported. Data from multiple systems cannot be

combined using standard control parameters.

ZONE: If data is being reported by zone, the zone number is shown.

Report Generated: The date on which the report was run.

ADDITIONAL ELEMENTS

Importance: Most reports show the relative importance assigned to the various service/report class performance goal types. See the CA MICS data dictionary entries for elements WLMGOLTP and WLMGOLRI. In addition, the Performance Manager Option defines importance level 6 and labels it Discretionary to give discretionary and 'other' work an importance distinct from the 'system' goal, which has a level of zero.

When reporting by WLM Report Class, this element may show the value Combined Work. This indicates that more than one workload type contributed to this report class period, thus making it a heterogeneous report class period.

Reporting Interval: The first row of some reports shows the length of time (duration) that contributed to the report.

% CHANGE: The last column on some reports, this field is the percent change for each report item (row) between the base period (first data column) and the last period (last data column). If data for only one period is reported, then dashes (---) are displayed. If the base period is zero, then N/A is printed to indicate that a percent change could not be calculated.

The following sections describe the reports produced by the WLM Goal Mode Service inquiries:

- 1 - WLM Service Definition Report
- 2 - WLM Service Delivered Summary, by Service Class
- 3 - WLM Service Delivered Summary, by Report Class
- 4 - WLM Service Delivered Summary Report

2.9.3.1 WLM Service Definition Report

The WLM Service Definitions report provides a simple, yet comprehensive, listing of the WLM Service Definitions in effect for the time interval selected. This listing provides an easy-to-follow method of understanding the definitions and the effect the definitions have on the workload performance. This information is helpful when trying to understand system performance behavior and in interpreting information from other WLM Goal Mode Service reports.

REPORT OPTIONS

Standard data selection options are supported as well as SYSID selection and external data selection. Additionally, input data may be further filtered by Sysplex or SVC (Service) Class. These options are discussed in Section 2.2.7. The reports may be ordered either by WLM Service Class (default), by WLM Importance, or by WLM Workload, as specified by a panel option.

REPORT PAGE HEADER FIELDS

The following fields are displayed at the beginning of each page of the report.

Sysplex:	The name of the SYSPLEX that the service definitions manage.
Service Definition:	The name of WLM Service Definition currently being reported on, from the CA MICS data element WLMSDFNM.
Description: (Svc. Def.)	The description of the WLM Service Definition named above, from the CA MICS data element WLMSDFDE.
Updated: (Svc. Def.)	The timestamp of the most recent update to the service definition, from CA MICS data element WLMSDFTS.
By:	The logon ID of the user who most recently updated the service definition, from CA MICS data element WLMSDUSR.
Service Policy:	The name of the WLM Service Policy in effect for this reporting interval, from CA MICS data element WLMPOLCY. A service policy may be thought of as a temporary, time-dependent

method of overriding goals for one or more service class periods. This might be done for different operational shifts or holidays, for example, when processing requirements might vary.

Description: The description of the WLM Service Definition (Svc. Policy) named above, from the CA MICS data element WLMPOLDE.

Updated: The timestamp of the most recent update to the (Svc. Policy) service definition, from CA MICS data element WLMPOLTS.

Ending Date: The date from the CA MICS Ending Timestamp of the input data.

Reporting Interval: The interval of the report, from the CA MICS common data element DURATION.

REPORT COLUMN HEADER FIELDS

Service Class: The WLM Service Class name, from CA MICS data element WLMCLASS.

Class Description: The description of the WLM Service Class named above, from the CA MICS data element WLMCLSDE.

Pd: The WLM Service Class period value, from CA MICS data element WLMCLSPD.

Imp: The WLM Importance value, from CA MICS data element WLMGOLRI.

Service Duration: The time interval, after which work will transition to a later WLM Service Class period, from CA MICS data element WLMPDDUR.

Workload: The name of the WLM Workload that the service class belongs to, from CA MICS data element WLMWRKLD.

Goal Type: The type of goal that the WLM Service Class period is managed to, from CA MICS data element WLMGOLTP. WLM Goal Types are either System Specified, Execution Velocity, Response Time or Discretionary, and are set by the user according to the type of workload processed.

%:	A ratio of the actual goal-defined service and the target service specified, expressed as a percentage and from CA MICS data element WLMGOLPC.
Goal (secs):	For response-time goals, the target response time, expressed typically in seconds and from CA MICS data element WLMGOLRT.
Resource Group:	The name of the WLM resource group in use by the service class, from CA MICS data element WLMGROUP. Resource groups are an optional method of limiting access to the processor. This could be used to prevent CPU-bound work from dominating a given service class, or to otherwise restrict work, perhaps under some sort of contractual arrangement. This option, if specified, is in effect for the entire Sysplex, not on a particular LPAR.
Protection CPU Stor:	Specifies whether work in a given service class is considered either CPU or Storage Critical. YES indicates a critical specification; a blank indicates no such specification.

CA											
CA MICS Performance Manager											
Listing of WLM Service Definitions											
Report Generated on: Monday, February 18, yyyy											
Sysplex: PLEXC1											
Service Definition: PLEXT1		Description: C1: 61,11,31,1,3,13				Updated: 24JAN08:14:41		By: WEGLESJ			
Service Policy: PLEXT1		Description: PLEXC1 - MERGED SYSPLEX				Updated: 24JAN08:14:41					
Ending Date: 02/16/08 Reporting Interval: 0:15:00											
Service	Class	Service	Class	Goal Resource	Protection	Unavail.					
Class	Description	Pd	Imp	Duration	Workload	Goal Type	%	(Secs)	Group	CPU	Stor
BATCHHOT	BATCH PRODUCTION HOT IMP 3 VEL30	1	3		BATCHPRD EXEC VELOCITY	30					RESPTIME
BATCHPRD	PRODUCTION BATCH - IMP 4 VEL 15	1	4		BATCHPRD EXEC VELOCITY	15					RESPTIME
BATCHTST	WLM TEST BATCH IMP 5 VEL 10	1	5		BATCHTST EXEC VELOCITY	10					RESPTIME
BATSTWLM	WLM TEST BATCH -3 PERIODS	1	3	2,500	BATCHTST EXEC VELOCITY	30					
DATABASE	DB2 & DATACOM MUFS - VELOCITY 20	1	3		DATABASE EXEC VELOCITY	20					RESPTIME
DATABPRD	PRODUCTION DB2/MUF IMP 2 VEL 45	1	2		DATABASE EXEC VELOCITY	45					RESPTIME
DATABSQA	QA DATACOM MUFS - VELOCITY 25	1	3		DATABASE EXEC VELOCITY	25					RESPTIME
DATBSHOT	HOT QA DATACOM TST IMP 2 VEL 30	1	2		DATABASE EXEC VELOCITY	30					RESPTIME
DDFDB2	DDF DB2 ENCLAVE	1	3	10,000	DDFDB2 RESP TIME PCTILE	80	5				RESPTIME
DISCRET	DISCRETIONARY SERVICE CLASS	1	6		DISCRET DISCRETIONARY						RESPTIME
ONLGOALP	CICS ONLINE PRODUCTION - GOAL	1	3		ONLINE RESP TIME PCTILE	90	0.2				RESPTIME
ONLPRDHT	HOT ONLINE PROD - IMP 2 VEL 25	1	2		ONLINE EXEC VELOCITY	25					RESPTIME
ONLPROD	PRODUCTION ONLINE- IMP 3 VEL 25	1	3		ONLINE EXEC VELOCITY	25					RESPTIME
ONLTEST	TEST ONLINE - VELOCITY 15	1	3		ONLINE EXEC VELOCITY	15					RESPTIME
OTHER	DEFAULT SERVICE CLS-DISCRETIONRY	1	6		OTHER DISCRETIONARY						RESPTIME
PCB_DFT	PROD CB DEFAULT ENCLAVES	1	3		WEB RESP TIME PCTILE	85	5				RESPTIME
PWEB_DFT	PROD WEB DEFAULT ENCLAVES	1	2	1,000	WEB RESP TIME PCTILE	85	5				RESPTIME
QA	QA TESTING - VELOCITY 25	1	3		QA EXEC VELOCITY	25					RESPTIME
QAHOT	DATACOM TESTING - CHARLES SCHWAB	1	1		STC EXEC VELOCITY	90					RESPTIME
ROSCOE	ROSCOE - VELOCITY 40- IMP 2	1	2		ROSCOE EXEC VELOCITY	40					
STCDB2	STC DB2 GROUP-VEL 40-IMP 2	1	2		STC EXEC VELOCITY	40					RESPTIME
STCDB2LO	STC DB2 GROUP-VEL 30-IMP 3	1	3		STC EXEC VELOCITY	30					RESPTIME
STCHOT	STC HOT-VEL45/IMP1 - NOT IN USE	1	1		STC EXEC VELOCITY	45					RESPTIME
STCMED	STC MEDIUM - VEL 30 IMP 2	1	2		STC EXEC VELOCITY	30					RESPTIME
STCNMD	STC NMD GROUP-VEL 25-IMP 2	1	2		STC EXEC VELOCITY	25					
STCPROD	STC PROD-MVS GROUP-VEL 40-IMP 1	1	1		STC EXEC VELOCITY	40					
STCTEST	STC TEST-USER STCS-VEL 20-IMP 3	1	3		STC EXEC VELOCITY	20					
SYSOTHER	UNCLASSIFIED WORK	1	0		SYSTEM SYSTEM SPECIFIED						RESPTIME
SYSSTC	STARTED TASK DEFAULT	1	0		SYSTEM SYSTEM SPECIFIED						
SYSTEM	HIGH PRIORITY SYSTEM WORK	1	0		SYSTEM SYSTEM SPECIFIED						RESPTIME
TSO	TSO - GOAL MODE - 3 PERIODS	1	2	2,500	TSO RESP TIME PCTILE	90	0.2				
TSOLOW	TSO - DB2 QUERY RELATED WORK	1	6		TSOLOW DISCRETIONARY				RG_NULL		RESPTIME
WEBSERV	TEMP SC FOR WEBSERVER STCS	1	3		WEB EXEC VELOCITY	50					RESPTIME
XSCVEL10	DUMMY SERVICE CLASS - 10%	1	3		SCVELDMY EXEC VELOCITY	10					RESPTIME
XSCVEL20	DUMMY SERVICE CLASS - 20%	1	3		SCVELDMY EXEC VELOCITY	20					RESPTIME
XSCVEL30	DUMMY SERVICE CLASS - 30%	1	3		SCVELDMY EXEC VELOCITY	30					RESPTIME
XSCVEL40	DUMMY SERVICE CLASS - 40%	1	3		SCVELDMY EXEC VELOCITY	40					RESPTIME
XSCVEL50	DUMMY SERVICE CLASS - 50%	1	3		SCVELDMY EXEC VELOCITY	50					RESPTIME
XSCVEL60	DUMMY SERVICE CLASS - 60%	1	3		SCVELDMY EXEC VELOCITY	60					RESPTIME
XSCVEL70	DUMMY SERVICE CLASS - 70%	1	3		SCVELDMY EXEC VELOCITY	70					RESPTIME
XSCVEL80	DUMMY SERVICE CLASS - 80%	1	3		SCVELDMY EXEC VELOCITY	80					RESPTIME
XSCVEL90	DUMMY SERVICE CLASS - 90%	1	3		SCVELDMY EXEC VELOCITY	90					RESPTIME

Figure 2-223. Sample WLM Service Definition Report

2.9.3.2 WLM Service Delivered Summary, by Service Class

This WLM Service Delivered report tracks information concerning transaction service delivered, resources consumed, and workload throughput and delays, at a monthly, weekly, daily, or hourly level.

Each logical page of the report presents the information for one service class and one service class period within one importance level for each single LPAR. This is also the order of the pages.

The report can also be summarized so that each page presents the information for one service class and one service class period within one importance level summarized for all LPARs in a sysplex.

When the zone option is specified for the inquiry, each page is expanded into a page for each zone in the incoming data that may have been filtered by data selection parameters.

When the hour option is specified for the inquiry, each page may be expanded into two or three pages to contain a column for each hour found in the incoming data that may have been filtered by data selection parameters.

Additionally, a CSV file can be optionally created for later use by Q&R to generate charts and graphs.

REPORT OPTIONS

Standard data selection options are supported as well as SYSID selection and external data selection. Additionally, input data may be further filtered by Sysplex or SVC (Service) Class. These options are discussed in Section 2.2.7.

DATA ELEMENTS

The elements reported are divided into the following categories:

- o Service Delivered
- o Resources Consumed
- o Throughput and Delays

SERVICE DELIVERED

This section presents measurement data about SRM transaction

volume and service received for transactions executing within the WLM Service Class Period being reported on.

Reporting Interval: The timespan, in seconds, covered by the report, from CA MICS common data element DURATION.

Total Transactions: The total number of SRM transactions that ended in the reporting interval, from CA MICS data element SECTORCP.

Transactions per Second: The rate of SRM transaction execution (ended transactions), per second, from CA MICS data element SECPSRCP.

Total SRM Svc Units (k): The total number of service units expended during the reporting interval, expressed in thousands. This value is derived from CA MICS data element SECTOSRV, divided by 1,000.

SRM Svc Units per Second: The rate of service units expended per second, expressed in thousands. This value is derived from CA MICS data element SECPSRV divided by 1,000.

Avg. Tx Queue Time (secs): The mean time transactions spent waiting for execution, expressed in seconds and from CA MICS data element SECAVQUE.

Avg. Tx Exec Times (secs): The mean time transactions spent in execution, expressed in seconds and from CA MICS data element SECAVXET.

RESOURCES CONSUMED

This section presents measurement data about processor, paging, and I/O resources used by transactions executing within the WLM Service Class Period being reported on.

Total CP Time (secs): The total time transactions executed on CPs (central processors), expressed in seconds and from CA MICS data element SECCPUTM.

zAAP on CP Time (secs): The total time the zAAP eligible Java work spent executing on CPs (central processors), expressed in seconds and from CA MICS data element SECZAPCT.

Total zAAP The total time transactions executed on zAAPs

Time (secs): (z/OS application assist processors), expressed in seconds and from CA MICS data element SECZAPTM.

zIIP on CP Time (secs): The total amount of TCB CPU time consumed within this service class period while executing zIIP eligible work on standard CP processors, from CA MICS data element SECSUPCT.

Total zIIP Time (secs): The total amount of normalized TCB CPU time executed on zIIP processors, expressed in seconds and from CA MICS data element SECSUPTM.

Total DASD I/O Count: The total number of non-paging I/O operations (SSCHs) done to DASD devices, expressed in thousands and from CA MICS data element SECDCNT.

\$SCOL

DASD I/Os per Second: The number of non-paging I/O operations (SSCHs) performed per second. This value is computed by dividing the Total DASD I/O Count defined above by the CA MICS data element DURATION.

DASD I/Os per CPU Second: The number of non-paging I/O operations (SSCHs) performed per CPU (TCB) second. This value is computed by dividing the Total DASD I/O Count defined above by the CA MICS data element SECCPUTM. This is a relative measure of workload I/O intensity and can be used to compare workloads within or across WLM service classes.

Avg. Page Ins per Second: The number of pages moved into central storage per second, from CA MICS data element SECPSPIN.

THROUGHPUT AND DELAYS

This section contains measurement data that tells how well transactions actually perform, including summary information about where processing delays occur. This provides the ability to do first-level problem analysis when service objectives are not being met.

Average MPL: The mean number of tasks active and executing within the WLM Service Class Period for the reporting interval, from CA MICS data element SECAVMPL.

Execution The relative measure of how fast transactions

Velocity:	are executing with the WLM Service Class Period, expressed as a percentage and from CA MICS data element SECEXVEL. Higher values indicate that transactions have experienced few delays; lower values indicate more delays. A zero value indicates that no processor resources were consumed during the reporting interval.
Performance Index:	A measure of how well transactions are performing with respect to their WLM-defined goals, from CA MICS data element SECPINDX. A value of 1 indicates that transactions are meeting their respective goals exactly. Values above 1 indicate that goals are being missed. Values below 1 indicate that goals are being exceeded.
CP Usage Delay Ratio:	The ratio of CPU Delay Samples to the sum of CPU Using and CPU Delay Samples. This is essentially the inverse of the CPU portion of the Execution Velocity calculation, and can be used to see when processor delays are occurring. Use this value in conjunction with the Execution Velocity and Performance Index to determine if service objectives are not being met and if processor delay is responsible.
zAAP Usage Delay Ratio:	The ratio of zAAP Delay Samples to the sum of zAAP Using and zAAP Delay Samples. This value can be used to see when zAAP delays are occurring. Use this value in conjunction with the Execution Velocity and Performance Index to determine if service objectives are not being met and if zAAP processor delay is responsible.
zIIP Usage Delay Ratio:	The ratio of zIIP Delay Samples to the sum of zIIP Using and zIIP Delay Samples. This value can be used to see when zIIP delays are occurring. Use this value in conjunction with the Execution Velocity and Performance Index to determine if service objectives are not being met and if zIIP processor delay is responsible.
Pct Processor Delays:	The sum of the following CA MICS data elements: SEDPCDCP (Pct Wait - CPU on CP) SEDPCZPD (Pct Wait - CPU on zAAP) SEDPCSUD (Pct Wait - CPU on zIIP)

This value represents the amount of processor delay as a function of all sampled delays.

Pct I/O
Related
Delays:

The sum of the following CA MICS data elements:
 SEDPCDSD (Pct Wait - DASD I/O Delay) ,
 SEDPCDHS
 (Pct Wait - Std Hiperspc Pg-In from Aux)
 SEDPCDES
 (Pct Wait - ESO Hiperspc Pg-In from Aux)
 SEDPCDVI (Pct Wait - VIO Page-In from Aux)

This value represents the amount of I/O delays as a function of all sampled delays.

Pct Storage
Delays:

The sum of the following CA MICS data elements:
 SEDPCDMP (Pct Wait - MPL Delay)
 SEDPCDSI (Pct Wait - Swap-In)
 SEDPCDPC (Pct Wait - CSA Page-In from Aux)
 SEDPCDPP (Pct Wait - Pvt Page-In from Aux)
 SEDPCDXM
 (Pct Wait - Cross Mem Page-In from Aux)

This value represents the amount of storage delays as a function of all sampled delays.

Pct Misc.
Delays:

The sum of the following CA MICS data elements:
 SEDPCCOD (Pct Wait - Resource Contention)
 SEDPCCRD (Pct Wait - Crypto Processors)
 SEDPCFQD (Pct Wait - Feature Queue)

This value represents the amount of other delays as a function of all sampled delays.

CA MICS Performance Manager Hourly WLM Service Delivered Summary For Sysplex: PLEXZA Sysid: CAZA Report Generated on: Monday, February 18, yyyy						
HOURLY-DATE (HR-MM/DD):	10-02/17	11-02/17	12-02/17	13-02/17	14-02/17	CHANGE %
Importance: System Tasks SVC Class: SYSSTC Period: 1						
Service Delivered						
Reporting Interval (secs)	3,600	3,600	2,432	3,600	3,600	0.00 %
Total Transactions	90	37	47	20	36	-60.00 %
Transactions per Second	0	0	0	0	0	-66.67 %
Total SRM Svc. Units (k)	24,614	20,279	12,222	17,995	16,370	-33.49 %
SRM Svc. Units per Sec.	6,838	5,633	5,024	4,999	4,547	-33.50 %
Avg. Tx. Queue Time (sec)	0.01	0.01	0.03	0.00	0.21	2000.00 %
Avg. Tx. Exec Time (sec)	1.79	1.02	25.04	0.77	1,118.98	2412.85 %
Resources Consumed						
Total CP Time (secs)	426.77	401.66	241.75	291.91	287.37	-32.66 %
- zAAP on CP Time (secs)	0.00	0.00	0.00	0.00	0.00	N/A
- zIIP on CP Time (secs)	0.00	0.00	0.00	0.00	0.00	N/A
Total zAAP Time (secs)	0.00	0.00	0.00	0.00	0.00	N/A
Total zIIP Time (secs)	0.00	0.00	0.00	0.00	0.00	N/A
Total DASD I/O Count (k)	618	270	280	224	194	-68.68 %
DASD I/Os per Second	172	75	115	62	54	-68.68 %
DASD I/Os per CPU Second	1,447.97	671.35	1,158.68	766.20	673.46	-53.49 %
Avg. Page Ins per Second	0	0	1	0	0	-85.71 %
Throughput and Delays						
Average MPL	65.27	68.86	63.75	64.61	64.79	-0.74 %
Execution Velocity	58.07	67.30	60.71	68.78	68.71	18.32 %
Performance Index	.	.	.	.	.	N/A
CP Usage Delay Ratio	0.41	0.32	0.38	0.31	0.31	-24.39 %
zAAP Usage Delay Ratio	0.00	0.00	0.00	0.00	0.00	N/A
zIIP Usage Delay Ratio	0.00	0.00	0.00	0.00	0.00	N/A
Pct Processor Delays	0.10	0.06	0.07	0.04	0.04	-60.00 %
Pct I/O Related Delays	0.18	0.02	0.10	0.02	0.01	-94.44 %
Pct Storage Delays	0.00	0.00	0.00	0.00	0.00	N/A
Pct Misc. Delays	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-224. Sample WLM Service Delivered Summary, by Service Class

2.9.3.3 WLM Service Delivered Summary, by Report Class

This WLM Service Delivered report tracks information concerning transaction service delivered, resources consumed, and workload throughput and delays, at a monthly, weekly, daily, or hourly level.

Each logical page of the report presents the information for one report class and report class period within one Importance level for each single LPAR. This is also the order of the pages.

The report can also be summarized so that each page presents the information for one report class and report class period within one Importance level summarized for all LPARs in a sysplex.

When the zone option is specified for the inquiry, each page is expanded into a page for each zone in the incoming data that may have been filtered by data selection parameters.

When the hour option is specified for the inquiry, each page may be expanded into two or three pages to contain a column for each hour found in the incoming data that may have been filtered by data selection parameters.

Additionally, a CSV file can be optionally created for later use by Q&R to generate charts and graphs.

REPORT OPTIONS

Standard data selection options are supported as well as SYSID selection and external data selection. Additionally, input data may be further filtered by Sysplex or RPT (Report) Class. These options are discussed in Section 2.2.7.

DATA ELEMENTS

The elements reported are divided into the following categories:

- o Service Delivered
- o Resources Consumed
- o Throughput and Delays

SERVICE DELIVERED

This section presents measurement data about SRM transaction volume and service received for transactions executing within the WLM Report Class Period being reported on.

Reporting Interval:	The timespan, in seconds, covered by the report, from CA MICS common data element DURATION.
Total Transactions:	The total number of SRM transactions that ended in the reporting interval, from CA MICS data element SECTORCP.
Transactions per Second:	The rate of SRM transaction execution (ended transactions), per second, from CA MICS data element SECPSRCP.
Total SRM Svc Units (k):	The total number of service units expended during the reporting interval, expressed in thousands. This value is derived from CA MICS data element SECTOSRV, divided by 1,000.
SRM Svc Units per Second:	The rate of service units expended per second, expressed in thousands. This value is derived from CA MICS data element SECPSSRV divided by 1,000.
Avg. Tx Queue Time (secs):	The mean time transactions spent waiting for execution, expressed in seconds and from CA MICS data element SECAVQUE.
Avg. Tx Exec Times (secs):	The mean time transactions spent in execution, expressed in seconds and from CA MICS data element SECAVXET.

RESOURCES CONSUMED

This section presents measurement data about processor, paging, and I/O resources used by transactions executing within the WLM Report Class Period being reported on.

Total CP Time (secs): The total time transactions executed on CPs (central processors), expressed in seconds and from CA MICS data element SECCPUTM.

zAAP on CP Time (secs): The total time the zAAP eligible Java work spent executing on CPs (central processors), expressed in seconds and from CA MICS data element SECZAPCT.

Total zAAP Time (secs): The total time transactions executed on zAAPs (z/OS application assist processors), expressed in seconds and from CA MICS data element SECZAPTM.

zIIP on CP Time (secs): The total amount of TCB CPU time consumed within this report class period while executing zIIP eligible work on standard CP processors, from CA MICS data element SECSUPCT.

Total zIIP Time (secs): The total amount of normalized TCB CPU time executed on zIIP processors, expressed in seconds and from CA MICS data element SECSUPTM.

Total DASD I/O Count: The total number of non-paging I/O operations (SSCHs) done to DASD devices, expressed in thousands and from CA MICS data element SECD CNT.

\$SCOL

DASD I/Os per Second: The number of non-paging I/O operations (SSCHs) performed per second. This value is computed by dividing the Total DASD I/O Count defined above by the CA MICS data element DURATION.

DASD I/Os per CPU Second: The number of non-paging I/O operations (SSCHs) performed per CPU (TCB) second. This value is computed by dividing the Total DASD I/O Count defined above by the CA MICS data element SECCPUTM. This is a relative measure of workload I/O intensity and can be used to compare workloads within or across WLM report classes.

Avg. Page Ins per Second: The number of pages moved into central storage per second, from CA MICS data element SECPSPIN.

THROUGHPUT AND DELAYS

This section contains measurement data concerning how well transactions actually perform, including summary information about where delays to processing occur. This provides the ability to do first-level problem analysis when service objectives are not being met.

Average MPL: The mean number of tasks active and executing within the WLM Report Class Period for the reporting interval, from CA MICS data element SECAVMPL.

Execution Velocity: The relative measure of how fast transactions are executing with the WLM Report Class Period, expressed as a percentage and from CA MICS data element SECEXVEL. Higher values indicate that transactions have experienced few delays; lower values indicate more delays. A zero value indicates that no processor resources were consumed during the reporting interval.

Performance Index: A measure of how well transactions are performing with respect to their WLM-defined goals, from CA MICS data element SECPINDX. A value of 1 indicates that transactions are meeting their respective goals exactly. Values above 1 indicate that goals are being missed. values below 1 indicate that goals are being exceeded.

CP Usage The ratio of CPU Delay Samples to the sum of
Delay Ratio: CPU Using and CPU Delay Samples. This is
essentially the inverse of the CPU portion of
the Execution Velocity calculation, and can be
used to see when processor delays are
occurring. Use this value in conjunction with
the Execution Velocity and Performance Index to
determine if service objectives are not being
met and if processor delay is responsible.

zAAP Usage The ratio of zAAP Delay Samples to the sum of
Delay Ratio: zAAP Using and zAAP Delay Samples. This value
can be used to see when zAAP delays are
occurring. Use this value in conjunction with
the Execution Velocity and Performance Index to
determine if service objectives are not being
met and if zAAP processor delay is responsible.

zIIP Usage The ratio of zIIP Delay Samples to the sum of
Delay Ratio: zIIP Using and zIIP Delay Samples. This value
can be used to see when zIIP delays are
occurring. Use this value in conjunction with
the Execution Velocity and Performance Index to
determine if service objectives are not being
met and if zIIP processor delay is responsible.

Pct Processor The sum of the following CA MICS data elements:
Delays: SEDPCDCP (Pct Wait - CPU on CP)
 SEDPCZPD (Pct Wait - CPU on zAAP)
 SEDPCSUD (Pct Wait - CPU on zIIP)

This value represents the amount of processor
delay as a function of all sampled delays.

Pct I/O
Related
Delays:

The sum of the following CA MICS data elements:
SEDPCDSD (Pct Wait - DASD I/O Delay) ,
SEDPCDHS
 (Pct Wait - Std Hiperspc Pg-In from Aux)
SEDPCDES
 (Pct Wait - ESO Hiperspc Pg-In from Aux)
SEDPCDVI (Pct Wait - VIO Page-In from Aux)

This value represents the amount of I/O delays
as a function of all sampled delays.

Pct Storage
Delays:

The sum of the following CA MICS data elements:
SEDPCDMP (Pct Wait - MPL Delay)
SEDPCDSI (Pct Wait - Swap-In)
SEDPCDPC (Pct Wait - CSA Page-In from Aux)
SEDPCDPP (Pct Wait - Pvt Page-In from Aux)
SEDPCDXM
 (Pct Wait - Cross Mem Page-In from Aux)

This value represents the amount of storage
delays as a function of all sampled delays.

Pct Misc.
Delays:

The sum of the following CA MICS data elements:
SEDPCCOD (Pct Wait - Resource Contention)
SEDPCCRD (Pct Wait - Crypto Processors)
SEDPCFQD (Pct Wait - Feature Queue)

This value represents the amount of other
delays as a function of all sampled delays.

CA MICS Performance Manager Hourly WLM Service Delivered Summary For Sysplex: PLEXZA Sysid: CAZA Report Generated on: Monday, February 18, yyyy						
HOURLY-DATE (HR-MM/DD):	10-02/17	11-02/17	12-02/17	13-02/17	14-02/17	CHANGE %
Importance: Imp. Level 1 RPT Class: STCPRD Period: 1						
Service Delivered						
Reporting Interval (secs)	3,600	3,600	2,432	3,600	3,600	0.00 %
Total Transactions	4	6	14	8	13	225.00 %
Transactions per Second	0	0	0	0	0	N/A
Total SRM Svc. Units (k)	3,725	3,558	2,800	9,467	3,375	-9.39 %
SRM Svc. Units per Sec.	1,035	988	1,151	2,630	937	-9.40 %
Avg. Tx. Queue Time (sec)	1.06	0.73	0.12	0.66	5.87	453.77 %
Avg. Tx. Exec Time (sec)	524.51	578.71	48.98	350.11	915.16	74.48 %
Resources Consumed						
Total CP Time (secs)	89.73	101.80	77.94	167.05	92.74	3.35 %
- zAAP on CP Time (secs)	0.00	0.00	0.00	0.00	0.00	N/A
- zIIP on CP Time (secs)	0.00	0.00	0.00	0.00	0.00	N/A
Total zAAP Time (secs)	0.00	0.00	0.00	0.00	0.00	N/A
Total zIIP Time (secs)	0.00	0.00	0.00	0.00	0.00	N/A
Total DASD I/O Count (k)	133	98	62	632	84	-36.43 %
DASD I/Os per Second	37	27	26	176	23	-36.44 %
DASD I/Os per CPU Second	1,479.34	963.47	800.52	3,783.05	909.90	-38.49 %
Avg. Page Ins per Second	0	0	0	0	0	-86.96 %
Throughput and Delays						
Average MPL	23.44	22.62	20.76	20.31	21.94	-6.40 %
Execution Velocity	35.34	41.12	31.98	57.45	62.36	76.46 %
Performance Index	1.13	0.97	1.25	0.70	0.64	-43.36 %
CP Usage Delay Ratio	0.64	0.58	0.68	0.43	0.37	-42.19 %
zAAP Usage Delay Ratio	0.00	0.00	0.00	0.00	0.00	N/A
zIIP Usage Delay Ratio	0.00	0.00	0.00	0.00	0.00	N/A
Pct Processor Delays	0.16	0.14	0.28	0.14	0.05	-68.75 %
Pct I/O Related Delays	0.04	0.03	0.04	0.17	0.02	-50.00 %
Pct Storage Delays	0.01	0.01	0.01	0.00	0.00	N/A
Pct Misc. Delays	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-225 Sample WLM Service Delivered Summary, by Report Class

2.9.3.4 WLM Service Delivered Summary Report

The WLM Service Delivered Summary report tracks information concerning CEC (Central Electronic Complex) and LP (Logical Partition) configuration and LP processor resource usage, as well as providing an overview of each workload's CP (Central Processor) and special purpose processor resource utilization. The report can be presented at a monthly, weekly, daily, or hourly level, and can be summarized either by WLM Importance, WLM Workload, WLM Service Class or WLM Report Class. Additionally, a CSV file can be optionally created for later use by Q&R to generate charts and graphs.

REPORT OPTIONS

Standard data selection options are supported as well as SYSID selection and external data selection. Additionally, input data may be further filtered by either SVC (Service) Class or RPT (Report) Class. These options are discussed in Section 2.2.7.

DATA ELEMENTS

The elements reported are grouped into the following sections:

- o CEC Configuration Summary
- o LP Configuration Summary
- o LP Processor Usage Summary
- o Workload Usage Summary

CEC CONFIGURATION SUMMARY

This section presents basic information concerning the CEC (Central Electronic Complex) where the LPAR being measured resides.

CEC Model Definition:	The model of the CEC under study, a concatenation of the CA MICS data elements CPUMOD (Processor Model Family) and CPCMODID (CPC Model Identifier).

CEC Serial Number:	The serial number of the CEC, taken from CA MICS data element PRSMCEC (PR/SM LPAR Central Complex).
Percent CEC Busy:	The total busy of all shared LPs in the CEC, expressed as a percentage and derived from the CA MICS data elements PRSMTTDT (Total Dispatch Time) and DURATION.
Shared CPs Defined in CEC:	The number of shared CPs defined within the CEC available for use in LPs, derived from CA MICS data elements PRSMACP (Avg CP Processors in Complex) and PRSMADP (Avg Dedicated Processors).
Avg zAAP Processors in CEC:	The average number of System z9 Application Assist Processors (zAAPs) assigned for use by the PR/SM LPAR feature, derived from CA MICS data element PRSMAZP (Avg zAAP Processors in Complex) and DURATION. This data element is valid only for System z9 and later servers.
Avg zIIP Processors in CEC:	The average number of System z9 Integrated Information Processors (zIIPs) assigned for use by the PR/SM LPAR feature, derived from CA MICS data element PRSMASP (Avg zIIP Processors in Complex) and DURATION. This data element is valid only for System z9 and later servers.

LP CONFIGURATION SUMMARY

This section presents measurement data about capacity sizing and processor resource usage within the LP being reported on.

Phys. CPU Adjust. Factor:	The number of CPU service units that can be expended in one clock second by CPs in this LP, from CA MICS data element CPUSUSEC (Physical CPU Adjustment Factor). This value is used to calculate the processor usage in MSUs, millions of service units per hour.
MSUs Avail. to Image:	The theoretical maximum capacity of the LP under study in MSUs per hour, from CA MICS data element CPUSUAVL and based on the CEC model and the number of CPs defined to the LP under measurement.

Weighted LP Cap. (MSUs):	The actual capacity of the LP under study in MSUs. This value is computed as a function of the CA MICS data elements CPUBPS (Partition Processor Weight) and PRSMTWGT (Sum of All Partitions Weights), and the MSUs Available to image value above. This is the amount of processor resources available for all CPs active within the LP when LP weighting is in effect.
Weighted LP Share Pct.:	The share of processor resources available for all active CPs in the LP under measurement, expressed as a percentage. This is the named result of the initial calculation of Weighted LP Cap. (MSUs) above and is effective once LP weight management is active.
No. of Available Procs.:	The maximum number of CPs and special purpose processors present in the configuration, from CA MICS data element CPUNOCPU (Max Number of Configured Processors).
Avg. CPs in Use:	The average number of CPs in use by the LP under study, from CA MICS data element CPUAVCPU (Avg Number of Assigned CP Processors).
Avg. zAAPs in Use:	The average number of zAAPs in use by the LP under study, from CA MICS data element CPUAVOZP (Avg Number of Online zAAP Processors).
Avg. zIIPs in Use:	The average number of zIIPs in use by the LP under study, from CA MICS data element CPUAVOSU (Avg Number of Online zIIP Processors).
CPs Required for LP Wgt.:	The minimum number of CPs required to sustain the weighted LP capacity. This value is computed as a function of the Weighted LP Share Percentage and the number of shared CPs defined in the CEC discussed previously. It is displayed as a decimal value. Round it upwards to the next integer and then compare the result against the Avg. CPs in Use value. If fewer CPs are in actual use, then potentially there will not be sufficient processor resources available to address the workload requirements.

LP Shr. Proc. Use Ratio: The ratio of LP processor usage as a function of z/OS busy time. This value is sometimes called "Short CPs" and is an indicator of work versus resources. Generally, a value greater than 1.25 is cause for concern and would indicate a potential performance impact.

Weighted CP Speed (MSUs): The effective speed of a single CP within the LP configuration. This is a function of the CPs Required for LP Wgt. calculation above and the Avg. CPs in Use value. This value becomes important when large processor tasks, such as CICS regions, are executing within the LP because the task execution can be frequently interrupted, causing poor or erratic response times within such regions.

LP PROCESSOR USAGE SUMMARY

This section contains measurement data concerning how processor resources are consumed within the LP. Information is included to help with planning for subcapacity license charging. Processor delays are presented.

Reporting Interval: The timespan, in seconds, covered by the report, from CA MICS common data element DURATION.

LP CP Busy Time (sec): The total processor busy time consumed by CPs assigned to this LP for the reporting interval, expressed in seconds. This value is from CA MICS data element CPUTOBTM (Total Time CP Processors Busy)

zAAP Eligible Time (sec):	The amount of TCB CPU time consumed within this service/report class period while executing zAAP eligible work on standard CP processors. This value is from CA MICS data element SECZAPCT (zAAP Eligible CPU Time on CP).
zIIP Eligible Time (sec):	The amount of TCB CPU time consumed within this service/report class period while executing zIIP eligible work on standard CP processors. This value is from CA MICS data element SECSUPCT (zIIP Eligible CPU Time on CP).
Percent LP CP Busy:	The total processor busy time consumed by CPs assigned to this LP for the reporting interval, expressed as a percentage. This value is from CA MICS data element CPUPCBSY (Pct CP Processors Busy).
LP zAAP Busy Time (sec):	The total processor busy time consumed by zAAPs assigned to this LP for the reporting interval, expressed in seconds. This value is from CA MICS data element CPUZPBTM (Total Time zAAP Processors Busy)
Percent LP zAAP Busy:	The total processor busy time consumed by zAAPs assigned to this LP for the reporting interval, expressed as a percentage. This value is from CA MICS data element CPUPCBZP (Pct zAAP Processors Busy).
LP zIIP Busy Time (sec):	The total processor busy time consumed by zIIPs assigned to this LP for the reporting interval, expressed in seconds. This value is from CA MICS data element CPUSUBTM (Total Time zIIP Processors Busy)
Percent LP zIIP Busy:	The total processor busy time consumed by zIIPs assigned to this LP for the reporting interval, expressed as a percentage. This value is from CA MICS data element CPUPCBSU (Pct zIIP Processors Busy).

Avg. MSUs Consumed:	The mean CP resources consumed, expressed in MSUs per hour. This value is derived using the LP CP Busy Time value described above and the Phys. CPU Adjust. Factor previously described. This value can be used in conjunction with the workload consumption values in the next section as a starting point for capacity planning. You can also use it to begin the subcapacity license charging process.
Uncaptured CP MSUs:	The CP resources that cannot be allocated to a given WLM defined workload and remain an overhead to the LP capacity.
Defined Capacity Limit:	The limit of LP capacity specified by the user in MSUs, from CA MICS data element CPUDEFLLM. If this value is specified, then LP capacity will be governed by use of the "four hour rolling average" of actual CP consumption; when the CP consumption exceeds this limit, then the LP is "soft capped."
Avg. 4 Hour Rolling MSUs:	The rolling average of LP consumption, in MSUs mentioned above and from CA MICS data element CPUAVLAC. This value can be tracked over time to help determine what an appropriate value for the Defined Capacity Limit should be if the user wants to implement subcapacity licensing charges.

Percent Ready Tasks Wait:	The time on average that tasks ready for dispatch must wait for an available CP, expressed as a percentage. This value is computed by first summing the appropriate CPURxx values and then dividing this sum by the SAMPLES value. Use this measure in conjunction with the LP Shr. Proc. Use Ratio to help identify intervals where a shortage of processor resources is affecting throughput.
CP Usage Delay Ratio:	The ratio of CPU Delay Samples to the sum of CPU Using and CPU Delay Samples. This is essentially the inverse of the CPU portion of the Execution Velocity calculation, and can be used to see when processor delays are occurring. Use this value in conjunction with the Execution Velocity and Performance Index to determine if service objectives are not being met and if processor delay is responsible.
zAAP Usage Delay Ratio:	The ratio of zAAP Delay Samples to the sum of zAAP Using and zAAP Delay Samples. This value can be used to see when zAAP delays are occurring. Use this value in conjunction with the Execution Velocity and Performance Index to determine if service objectives are not being met and if zAAP processor delay is responsible.
zIIP Usage Delay Ratio:	The ratio of zIIP Delay Samples to the sum of zIIP Using and zIIP Delay Samples. This value can be used to understand when zIIP delays are occurring. Use this value in conjunction with the Execution Velocity and Performance Index to determine if service objectives are not being met and if zIIP processor delay is responsible.

CA MICS Performance Manager Hourly LPAR Service Delivered Summary Sysid: CAZA Report Generated on: Monday, February 18, yyyy						
HOURLY-DATE (HR-MM/DD):	10-02/16	11-02/16	12-02/16	13-02/16	14-02/16	CHANGE %
CEC Configuration Summary						
CEC Model Definition	2084-312	2084-312	2084-312	2084-312	2084-312	
CEC Serial Number	00851	00851	00851	00851	00851	
Percent CEC Busy	95.86	96.59	96.47	96.10	95.85	-0.01 %
Shared CPs Defined in CEC	12	12	12	12	12	---
Avg zAAP Proc. in CEC	0	0	0	0	0	N/A
Avg zIIP Proc. in CEC	0	0	0	0	0	N/A
LP Configuration Summary						
Phys. CPU Adjust. Factor	14,350	14,350	14,350	14,350	14,350	---
MSUs Available to Image	207	207	207	207	207	---
Weighted LP Cap. (MSUs)	36	36	36	36	36	---
Weighted LP Share Pct.	17.2	17.2	17.2	17.2	17.2	---
No. of Available Procs.	4	4	4	4	4	---
Avg. CPs in Use by LP	3.00	3.00	3.00	3.00	3.00	---
Avg. zAAPs in Use by LP	0.00	0.00	0.00	0.00	0.00	N/A
Avg. zIIPs in Use by LP	0.00	0.00	0.00	0.00	0.00	N/A
CPs Required for LP Wgt.	2.07	2.07	2.07	2.07	2.07	---
LP Shr. Proc. Use Ratio	1.00	1.00	1.00	1.00	1.00	---
Weighted CP Speed (MSUs)	36	36	36	36	36	---
LP Processor Usage Summary						
Reporting Interval (sec)	3,600	3,600	3,600	3,600	3,600	---
LP CP Busy Time (sec)	5,837	7,323	5,689	5,392	5,252	-10.02 %
- zAAP Eligible Time(sec)	432	564	421	389	234	-3.02 %
- zIIP Eligible Time(sec)	432	564	421	389	234	-3.02 %
Percent LP CP Busy	54	68	53	50	49	-10.01 %
LP zAAP Busy Time (sec)	0	0	0	0	0	N/A
Percent LP zAAP Busy	0	0	0	0	0	N/A
LP zIIP Busy Time (sec)	0	0	0	0	0	N/A
Percent LP zIIP Busy	0	0	0	0	0	N/A
Avg. MSUs Consumed	84	105	82	77	75	-10.02 %
- Uncaptured CP MSUs	12	15	13	13	13	8.33 %
Defined Capacity Limit	0	0	0	0	0	N/A
Avg. 4 Hour Rolling MSUs	81	83	84	84	86	5.56 %
Percent Ready Tasks Wait	30	51	33	32	33	11.84 %
CPU Usage Delay Ratio	0.80	0.70	0.40	0.60	0.70	-12.50 %
zAAP Usage Delay Ratio	0.00	0.00	0.00	0.00	0.00	N/A

Figure 2-226 Sample LPAR Service Production Summary Report

2.10 Product Measured Usage Analysis

This section introduces the concepts of IBM's z/OS subcapacity pricing option to form a background for discussing Product Measured Usage analysis. Subcapacity pricing is IBM's methodology for charging for certain software based on the concept of the highest observed four-hour rolling average utilization of the LPARs in which the products run rather than based on the MSU rating of the machine.

Section 2.10.4, Subcapacity Pricing Concepts, defines the terms and explains the basic concepts of subcapacity pricing. Depending on your familiarity with these terms and concepts, you may find it useful to read this section before proceeding with Section 2.10.2.

The following sections describe the operation of the Product Measured Usage Analysis inquiry and how to use the reports produced:

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output
- 4 - Subcapacity Pricing Concepts
- 5 - Reporting Requirements

2.10.1 Functional Description

The Product Measured Usage Analysis queries allow you to produce a software inventory by vendor and LPAR, and track product LPAR utilization. It will also provide a mechanism to report on the combined hourly system usage for each LPAR in which a product ran, and contrast the system usage against the combined hourly product usage for each of the same LPARs.

The LPAR Product Resource Consumption query uses data from the Measured Usage Global (BATMUG) File. The Combined Hourly System Usage by Product query processes both the BATMUG file and the CPU Activity (HARCPU) File. In the later case, both the CA MICS Batch and Operations Analyzer (SMF) DIA and the CA MICS Hardware and SCP Analyzer (RMF) DIA have to be installed in the same unit.

The range of data that can be reported from a particular CA MICS database depends on the timespans that were defined for these files and the number of cycles defined in each timespan. If you are not familiar with the generation options for the database unit containing the Batch Activity or the Hardware Utilization Information Area, ask your system administrator for help.

2.10.2 Usage Guidelines

Subcapacity pricing is IBM's charging methodology for certain software based on the utilization capacity of the LPARs in which the products run. This pricing method offers LPAR-based pricing for subcapacity eligible software products, based on the highest four-hour rolling average utilization of the LPARs where the product ran. There are many advantages and disadvantages to moving to a subcapacity pricing method, so careful planning and analysis is important. You will need to understand the prerequisites for subcapacity pricing, create a software inventory, create a capacity and growth plan, and determine your current capacity and your system usage in contrast to your product usage across all LPARs in which a product ran.

The CA MICS Performance Manager, IBM License Manager (ILM) Analysis queries, and the CA MICS Capacity Planner can determine your current capacity and usage for each LPAR and help you develop a capacity plan for the future. See the CA MICS Capacity Planner guide for more information on capacity planning for z/OS.

The Product Measured Usage queries can help you analyze product resource consumption, create a software inventory, and report the combined system usage in contrast to the combined product usage across all LPARs.

The following sections present guidelines to consider when using the available queries to report on product measured usage.

- 1 - LPAR Product Resource Consumption
- 2 - Combined Hourly System Usage by Product

2.10.2.1 LPAR Product Resource Consumption

The LPAR Product Measured Usage query analyzes information in the Measured Usage Global (BATMUG) file. It can help you track resource consumption and produce a software inventory by LPAR for any products reported with the SMF type 89 (Usage Data) records. Resource consumption analysis for performance or workload shifting is best done at the DAYS-or-DAYS-by-Hour summarization level.

To produce a general software inventory, it is best to execute the LPAR Product Measured Usage report against the BATMUG file at the MONTHS timespan. This report will summarize the entire month's product usage by LPAR by product into a one-line-per-product-by-month report.

When generating software inventory for the purpose of subcapacity pricing, the report should be run against the DAYS timespan, producing a report summarizing product usage by day by product. The problem with generating this report for this purpose from the month's timespan is that your CA MICS month's file will probably be defined from the first day of one month to the last day of the next month. IBM defines its billing cycle as the second day of one month to the first day of the next month. Therefore, you will either be including the first day of the prior month or excluding the first day of the next month.

To help you understand what is in the BATMUG file, the four CA MICS elements that make up a product follow in their key file sequence order:

- MUPROD - Product Name
- MUVER - Product Version
- MUPRODID - Product ID
- MUQUAL - Product Qualifier

By default, the report presents product usage summarized at the product ID level. Optionally, the report can be reported at the product qualifier level. When reporting at the product qualifier level you will see all instances of an individual product's execution. For example, when reporting at the product qualifier level, your report will show all instances of IMS/ESA V9R1 batch and started tasks. Conversely, when reporting at the product ID level, all IMS/ESA V9R1 usage will be summarized to the product ID level.

The the input data can be subset by specifying a vendor in the Data Selection section of the query composition panel.

For example, if you only wanted to report on product supplied by IBM, you would type 'IBM CORP' in the Vendor field. The data will then be subset prior to reporting.

To further subset the input data beyond the data selection options available on the query composition panel, you can specify a global data selection exit by coding a SELECT macro under the Extended Options panel. If you enter Y (Yes) in the Specify global data selection exit field, the Global Data Selection Panel, which is shown in Figure 2-227, is displayed. This panel allows you to enter SAS statements that will perform special data selection functions.

The SELECT MACRO can be used for special selection code based on any data element in the BATMUG file. A typical use may be to limit data analysis by excluding a specific product from analysis. The following example will exclude all releases of z/OS from the analysis.

```
/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step: Combined Hourly System Usage Analysis |
| Line Cmds: I Insert  D Delete  R Repeat  M Move   C Copy |
| Cmd  SAS Statements for: Specifying global data selection |
| - ----- |
| - %MACRO SELECT; |
| - IF MUPROD NE 'z/OS'; ----- |
| - ----- |
| - %MEND SELECT; |
| ***** BOTTOM OF DATA ***** |
|-----/
```

Figure 2-227. Global Data Selection Exit Panel

The generated report shows product consumption and maximum product consumption in CPU time, service units (SUs), and millions of service units (MSUs). It calculates a product's MSU consumption rate based on observed activity.

The Product Measured Usage inquiry lets you analyze the following items:

- o Product Resource Consumption in CPU Time
- o Product Resource Consumption in SUs
- o Product Resource Consumption in MSUs
- o Product Resource Consumption as an MSU Rate

2.10.2.2 Combined Hourly System Usage by Product

The Combined Hourly System Usage by Product reports analyze information from the Measured Usage Global (BATMUG) file, merged with the Hardware Utilization (HARCPU) file. The first two BATMUG Frequency Reports will provide a software inventory of any products reported with the SMF type 89 (Usage Data) records along with a list of vendors represented in the data. The last BATMUG Frequency Report provides a matrix of products versus the LPARs on which they ran. Use the information provided in these reports to subset the input data, reducing the size and scope of the reports.

To further subset the input data beyond the data selection options available on the query composition panel, you can specify a global data selection exit by coding a SELECT macro under the Extended Options panel. If you enter Y (Yes) in the Specify global data selection exit field, the Global Data Selection Exit Panel, which is shown in Figure 2-228, is displayed. This panel lets you enter SAS statements that will perform special data selection functions.

The SELECT MACRO can be used for special selection code based on any data element in the BATMUG file. A typical use may be to limit data analysis by excluding a specific product from analysis. The following example will exclude all releases of z/OS from the analysis.

```
/----- SAS Statements -----\
| Command ==>                               Scroll ==> CSR |
| Inquiry Step: Combined Hourly System Usage Analysis |
| Line Cmds: I Insert D Delete R Repeat M Move C Copy |
| Cmd SAS Statements for: Specifying global data selection |
| - ----- |
| - %MACRO SELECT; |
| - IF MUPROD NE 'z/OS'; ----- |
| - ----- |
| - %MEND SELECT; |
| ***** BOTTOM OF DATA ***** |
\-----/
```

Figure 2-228. Global Data Selection Exit Panel

Because IBM defines its billing cycle as the second day of one month to the first day of the next month, when generating a software inventory for the purpose of subcapacity pricing, the query must be set up to report on data from that time period only.

The All-Hour report will do the following:

- o Determine the LPARs on which each product executed.
- o Calculate and sum the product's MSU rate for each hour and LPAR in which it ran.
- o Sum the LPARs' four-hour rolling averages across all LPARs on which the product ran.
- o Generate a report section for each hour the product ran.
- o Display the combined systems' four-hour rolling average along with the combined products' MSU rate.

Subcapacity pricing is based on the concept of the highest observed rolling four-hour average utilization of the LPAR or LPARs where a product runs. The simultaneous combined utilization of these LPARs is determined for each hour in the reporting period and the highest observed combined utilization is used as the basis for pricing the product.

The High-Low Combined System Usage report is simply a subset of the All-Hour Combined System Usage Report. It is designed to display the four highest observed combined utilization hours and the four lowest for each hour the product ran. Using this information you can begin to study the possible effects of workload shifting in respect to IBM's subcapacity pricing methodology. There are many advantages and disadvantages to moving to a subcapacity pricing method, so careful planning and analysis is important.

2.10.3 Standard Output

COMMON REPORT ELEMENTS

Each report has several lines of title information that give the name of the report, the dates being reported, the data summarization level, the product summary level, and the start and end times of the period covered by the report. The following elements are in the standard title part of each report:

Report Name:	The name of the report, centered in the first line of the heading. It cannot be modified.
SYSPLEX:	The name of the sysplex from which the data originated.
SYSID or System:	The CA MICS logical SYSID being reported. Data from multiple systems cannot be combined using standard control parameters.
DATE:	The date represents the day, week, or month for the data being displayed on that page of the report.
ZONE:	If data is being reported by zone, the zone number is shown.
REPORT SUMMARY:	Indicates the summary level to which the reporting data has been summarized and displayed.
REPORT RANGE:	Indicates the entire range of data that is being reported.
PRODUCT SUMMARY:	Indicates to which level of product the report data has been summarized. There are two levels at which this query can be summarized: Product Name Version ID Qualifier or Product Name Version ID

REPORT OPTIONS

The CA MICS Performance Manager reports are initially delivered with defaults. These defaults will change based on the settings of the last execution of the query. For example, the LPAR Product Consumption report is delivered with a hard-coded default value of DAYS for the timespan. If

the user changes the timespan so that report will run at the WEEKS timespan, then the next execution of this report will default to the WEEKS timespan.

All reports support standard time selection options, as well as SYSID selection and external data selection. They also support subsetting the input data based on the Vendor option. When a vendor is specified, only that data will be selected for reporting. These options are controlled by the control options on the query composition panel.

The following analysis reports are discussed in more detail:

- 1 - LPAR Product Consumption Report
- 2 - Combined Hourly System Usage Reports

2.10.3.1 LPAR Product Consumption Report

The LPAR Product Resource Consumption query analyzes BATMUG data collected from z/OS workloads. It lets you produce a software inventory by vendor and LPAR. To produce a general software inventory by LPAR, it is best to execute the LPAR Product Resource Consumption query against the MONTHS timespan. The report summarizes the entire month's product usage by LPAR by product into a one-line-per-product-by-month report.

To study product usage for the purpose of performance analysis or workload shifting, this query must be executed at the DAYS-by-Hour summarization level. This query will produce a report that tracks product LPAR utilization in CPU Time, Service Units(SUs), Millions of Service Units(MSUs), and as an MSU Hourly Rate.

This report can be produced at multiple summarization levels. Dates and titles will be adjusted appropriately based on the timespan and summarization level chosen for the execution. After data has been collected for a time period, this summarization capability can be used to "drill down" to a specific time frame.

By default, the report presents products for all vendors. Optionally, by specifying a vendor on the composition panel, you can produce a report for only the desired vendor. To obtain the specific vendor name to subset on, you can first run the query against all vendor products. View the output and note the specific vendor name you want to further refine your query against. Set the vendor option field to the specific vendor you are interested in, and resubmit your query.

By default, the report presents product usage summarized at the product-ID level. Optionally, by modifying an execution panel field, the report will be summarized at the product-qualifier level. Consider an example:

- o When reporting at the product-qualifier level, you can separate resource collection for different iterations of the same product (MUPROD) running on a single system. For example, your report will show all instances of IMS/ESA V9R1 batch and started tasks separately.
- o Conversely, when reporting at the product-ID level, all IMS/ESA V9R1 usage would be summarized at the product-ID level.

Additionally, the LPAR Product Resource Consumption query provides the ability to produce the BATMUG Frequency Reports at either the product-ID or the product-qualifier level. See section 2.10.3.2.1 below for details.

A CSV file can be optionally created for later use by Q&R Workstation to generate charts and graphs.

REPORT ELEMENTS

Report Label	File	Element
-----	----	-----
Product Name	BATMUG	MUGPROD
Product Version	BATMUG	MUGVER
Product ID	BATMUG	MUGID
Product Qualifier	BATMUG	MUGQUAL
Owner/Vendor Name	BATMUG	MUOWNER
Product Usage Data Duration	BATMUG	MUGDUR
Total Product CPU Time	BATMUG	MUGCPUTM
Max Product CPU Time	BATMUG	MUGMXCPU

LPAR PRODUCT CONSUMPTION REPORT

To help you better understand the standard output produced by this query, definitions and calculations involved in producing the displayed variables follow:

TOTAL PRODUCT USAGE:	The total CPU time consumed by the product over the reporting interval.
MAX PRODUCT USAGE:	The maximum CPU time consumed by the product over the reporting interval. If the report is by hour, this value will represent the maximum CPU time consumed for that hour.
TOTAL PRODUCT USAGE IN SUs:	The number of service units consumed by the product over the reporting interval. $\text{PRODSUTO} = \text{MUGCPUTM} * \text{SUCONFA}$ MUGCPUTM is the sum of TCB and SRB CPU time, in hundredths of a second, as recorded by the Measured Usage License Charging (MULC) subsystem. SUCONFA is the physical CPU adjustment factor

used to convert CPU times to SUs.

The CPU adjustment factor is calculated by multiplying the CPU Capability Adj. Factor for N CPUs (CPCMAFxx), the CPU Capability (CPCCPUCB), and the Configured CPU Count (CPCCFGCT). All are from the SMF type 89 (Usage Data) record.

MAX
PRODUCT
USAGE IN
SUs: The maximum number of service units consumed by the product over the reporting interval. If reporting by hour, this value will represent the maximum SUs consumed for that hour.

TOTAL
PRODUCT
USAGE IN
MSUs: The number of service units (expressed in MSUs) consumed by the product for the reporting interval.

MAX
PRODUCT
USAGE IN
MSUs: The number of service units (expressed in MSUs) consumed by the product for the reporting interval. If reporting by the hour, this value will represent the maximum MSUs consumed for that hour.

PRODUCT
HOURLY
MSU RATE: Contains the product's hourly MSU consumption rate.

$$\text{PRDMSUR8} = \frac{\text{MUGCPUTM} * \text{SUCONFA} * 3600}{\text{DURATION} * 1000000}$$

MUGCPUTM is the sum of TCB and SRB CPU time, in hundredths of a second, as recorded by the Measured Usage License Charging (MULC) subsystem.

SUCONFA is the physical CPU adjustment factor used to convert CPU times to SUs.

The CPU adjustment factor is calculated by multiplying the CPU Capability Adj. Factor for N CPUs (CPCMAFxx), the CPU Capability (CPCCPUCB), and the Configured CPU Count (CPCCFGCT). All are from the SMF type 89 (Usage Data) record.

DURATION is the length of the SMF recording interval.

L P A R P r o d u c t C o n s u m p t i o n R e p o r t											Page 001
Sysplex Name: PLEXC1 SYSID: SA31 DAY: yy-mm-dd				Report Summary: Information Summarized By Hour Within Day							
				Report Range: 18AUGyy:10:59:37 thru 19AUGyy:12:59:37							
				Product Summary: Product Name Version ID							
Hour: 10				Total Product CPU Time (TCB+SRB)	Max Product CPU Time (TCB+SRB)	Total Product Usage (SUs)	Max Product Usage (SUs)	Total Product Usage (MSUs)	Max Product Usage (MSUs)	Product Hourly MSU Rate	
Product Name	Product Version	Product Id	Owner/Vendor Name								
z/OS	01.08.00	5645-001	IBM CORP	0:00:54.1	0:00:30.3	708.7k	397.1k	0.71	0.40	0.71	
CICS	V2 R2.0	5697-E93	IBM CORP	0:00:00.8	0:00:00.4	11.0k	5.5k	0.01	0.01	0.01	
CICS	V2 R3.0	5697-E93	IBM CORP	0:00:01.6	0:00:00.8	20.5k	10.3k	0.02	0.01	0.02	
CICS	V3 R1.0	5655-M15	IBM CORP	0:00:03.6	0:00:01.8	47.5k	23.9k	0.05	0.02	0.05	
CICS	V3 R2.0	5655-M15	IBM CORP	0:00:01.9	0:00:01.0	25.1k	12.7k	0.03	0.01	0.03	
DB2	7.1.0	5675-DB2	IBM CORP	0:00:05.3	0:00:02.7	69.2k	35.1k	0.07	0.04	0.07	
DB2	8.1.0	5625-DB2	IBM CORP	0:00:06.0	0:00:03.0	78.1k	39.5k	0.08	0.04	0.08	
DB2	9.1.0	5635-DB2	IBM CORP	0:00:01.0	0:00:00.6	13.5k	7.3k	0.01	0.01	0.01	
DB2	9.1.0	5635-DB2	IBM CORP	0:00:01.0	0:00:00.6	13.5k	7.3k	0.01	0.01	0.01	
IMS/ESA	V9R1	5655-J38	IBM CORP	0:00:00.7	0:00:00.4	9.3k	4.8k	0.01	0.00	0.01	
MQM MVS/ESA	V6	5655-L82	IBM CORP	0:00:13.7	0:00:03.2	178.9k	42.0k	0.18	0.04	0.18	
Hour: 11				Total Product CPU Time (TCB+SRB)	Max Product CPU Time (TCB+SRB)	Total Product Usage (SUs)	Max Product Usage (SUs)	Total Product Usage (MSUs)	Max Product Usage (MSUs)	Product Hourly MSU Rate	
Product Name	Product Version	Product Id	Owner/Vendor Name								
z/OS	01.08.00	5645-001	IBM CORP	0:00:48.2	0:00:25.5	630.5k	333.8k	0.63	0.33	0.63	
CICS	V2 R2.0	5697-E93	IBM CORP	0:00:00.8	0:00:00.4	11.1k	5.6k	0.01	0.01	0.01	
CICS	V2 R3.0	5697-E93	IBM CORP	0:00:01.6	0:00:00.8	20.9k	10.6k	0.02	0.01	0.02	
CICS	V3 R1.0	5655-M15	IBM CORP	0:00:03.7	0:00:01.8	48.2k	24.2k	0.05	0.02	0.05	
CICS	V3 R2.0	5655-M15	IBM CORP	0:00:02.0	0:00:01.0	26.0k	13.1k	0.03	0.01	0.03	
DB2	7.1.0	5675-DB2	IBM CORP	0:00:05.3	0:00:02.7	68.8k	34.8k	0.07	0.03	0.07	
DB2	8.1.0	5625-DB2	IBM CORP	0:00:05.7	0:00:03.0	75.0k	39.0k	0.07	0.04	0.07	
DB2	9.1.0	5635-DB2	IBM CORP	0:00:01.0	0:00:00.5	13.3k	7.2k	0.01	0.01	0.01	
IMS/ESA	V9R1	5655-J38	IBM CORP	0:00:00.7	0:00:00.3	8.9k	4.6k	0.01	0.00	0.01	
MQM MVS/ESA	V6	5655-L82	IBM CORP	0:00:15.4	0:00:03.9	201.4k	50.9k	0.20	0.05	0.20	
Hour: 12				Total Product CPU Time (TCB+SRB)	Max Product CPU Time (TCB+SRB)	Total Product Usage (SUs)	Max Product Usage (SUs)	Total Product Usage (MSUs)	Max Product Usage (MSUs)	Product Hourly MSU Rate	
Product Name	Product Version	Product Id	Owner/Vendor Name								
z/OS	01.08.00	5645-001	IBM CORP	0:06:02.3	0:05:07.5	4.7m	4.0m	4.74	4.02	4.74	
CICS	V2 R2.0	5697-E93	IBM CORP	0:00:00.8	0:00:00.4	10.9k	5.5k	0.01	0.01	0.01	
CICS	V2 R3.0	5697-E93	IBM CORP	0:00:01.6	0:00:00.8	20.4k	10.3k	0.02	0.01	0.02	
CICS	V3 R1.0	5655-M15	IBM CORP	0:00:03.6	0:00:01.8	47.0k	23.9k	0.05	0.02	0.05	
CICS	V3 R2.0	5655-M15	IBM CORP	0:00:02.0	0:00:01.0	25.8k	13.1k	0.03	0.01	0.03	
DB2	7.1.0	5675-DB2	IBM CORP	0:00:05.2	0:00:02.7	68.7k	35.6k	0.07	0.04	0.07	
DB2	8.1.0	5625-DB2	IBM CORP	0:00:05.8	0:00:03.0	75.4k	38.9k	0.08	0.04	0.08	
DB2	9.1.0	5635-DB2	IBM CORP	0:00:01.0	0:00:00.5	13.7k	7.2k	0.01	0.01	0.01	
IMS/ESA	V9R1	5655-J38	IBM CORP	0:00:00.7	0:00:00.3	8.9k	4.4k	0.01	0.00	0.01	

Figure 2-229. Sample LPAR Product Consumption Report

2.10.3.2 Hourly System Usage by Product Reports

These reports are designed to show at a glance which products ran on which systems, the total system usage for each system the product executed on, and the total system resources consumed by the product on each system. By comparing these numbers, you can get an idea of what percentage of the system is used by each product.

This section describes elements common to these reports. Options that are unique to a specific report are documented with that report.

COMMON REPORT ELEMENTS

Report Label	File	Element
-----	----	-----
Product Name	BATMUG	MUGPROD
Product Version	BATMUG	MUGVER
Product ID	BATMUG	MUGID
Product Qualifier	BATMUG	MUGQUAL
Product Usage Data Duration	BATMUG	MUGDUR
Total Product CPU Time	BATMUG	MUGCPUTM

PRODUCT: is a concatenation of the following three CA MICS data elements, which form a unique entity for reporting purposes.

MUPROD is the name of the software product participating in Measured Usage License Charging (MULC).

MUVER is the product version identifier for a software product participating in Measured Usage License Charging.

MUPRODID is the name of the product identifier for a software product participating in Measured Usage License Charging. IBM uses the product information department (PID) numbers to populate product identifier.

The following sections describe the report elements and the report formats of the BATMUG Frequency Reports, the High-Low Hourly Usage Report, and the All-Hour Hourly Usage Report.

- 1 - BATMUG Frequency Reports
- 2 - High-Low Hourly Usage Report
- 3 - All-Hour Hourly Usage Report

2.10.3.2.1 BATMUG Frequency Reports

The Product Measured Usage Analysis inquiries give you the option to generate frequency reports against the BATMUG file at the report-summary level:

- o Since the Hourly System Usage by Product inquiry is designed to report at the product-ID level, the corresponding frequency reports will report on the BATMUG information summarized at the product-ID level.
- o The LPAR Product Resource Consumption inquiry documented in section 2.10.3.1 provides the ability to produce the BATMUG Frequency Reports at either the product-ID or the product-qualifier level.

The BATMUG Frequency Reports provide a software inventory of any products reported with the SMF type 89 (Usage Data) records along with a list of vendors represented in the data. The last BATMUG Frequency Report provides a matrix of products by the LPARs on which they ran.

To help you better understand the standard output produced by this inquiry, definitions of variables displayed in the produced reports follow:

OWNER/ VENDOR NAME:	is the name of the product owner or vendor for a software product participating in Measured Usage License Charging (MULC).
PRODUCT:	is a concatenation of the first three CA MICS data elements, which form a unique entity for reporting purposes. When reporting at the product-qualifier level, MUQUAL is concatenated after MUPRODID.
MUPROD	is the name of the software product participating in Measured Usage License Charging (MULC).
MUVER	is the product version identifier for a software product participating in Measured Usage License Charging.
MUPRODID	is the name of the product identifier for a software product participating in Measured Usage License Charging. IBM uses the product information department (PID) numbers to populate

product identifier.

MUQUAL contains the product qualifier for a software product participating in Measured Usage License Charging.

Frequency Report on MUOWNER:Vendor/Owner 12:58 Wednesday, February 13, yyyy

The FREQ Procedure

Owner/Vendor Name

MUOWNER	Frequency
CA	20180
IBM CORP	1196

Frequency Report on Product:Product Name|Version|Product ID
12:58 Wednesday, February 13, yyyy

The FREQ Procedure

MUPROD|MUSER|MUPRODID

PRODUCT	Frequency
z/OS 01.08.00 5645-001	168
z/OS 01.09.00 5645-001	48
CICS V1 R2.0 5655-147	24
CICS V2 R2.0 5697-E93	92
CICS V2 R3.0 5697-E93	63
CICS V3 R1.0 5655-M15	58
CICS V3 R1.0 5697-E93	4
CICS V3 R2.0 5655-M15	36
CICS V4 R1.0 5655-018	48
DB2 7.1.0 5675-DB2	144
DB2 8.1.0 5625-DB2	96
DB2 9.1.0 5635-DB2	72
IMS/ESA V0R1 5635-A01	5
IMS/ESA V8R1 5655-C56	31
IMS/ESA V9R0 5655-J38	16
IMS/ESA V9R1 5655-J38	56
MQM MVS/ESA V6 5655-L82	235

Frequency Report on PRODUCT*SYSID 12:58 Wednesday, February 13, yyyy										
The FREQ Procedure										
Table of PRODUCT by SYSID										
PRODUCT(MUPROD MUPER MUPRODID)	SYSID(System Identifier)									
Frequency	LP11	LP31	LPD1	LP11	LPT1	LP55	LP09	LP21	LP61	Total
z/OS 01.08.00 5645-001	24	0	24	24	24	24	24	24	0	168
z/OS 01.09.00 5645-001	0	24	0	0	0	0	0	0	24	48
CICS V1 R2.0 5655-147	0	0	0	0	0	0	24	0	0	24
CICS V2 R2.0 5697-E93	17	3	0	24	24	0	0	24	0	92
CICS V2 R3.0 5697-E93	15	23	1	0	0	0	0	24	0	63
CICS V3 R1.0 5655-M15	17	23	5	0	0	0	0	0	13	58
CICS V3 R1.0 5697-E93	4	0	0	0	0	0	0	0	0	4
CICS V3 R2.0 5655-M15	12	24	0	0	0	0	0	0	0	36
CICS V4 R1.0 5655-018	0	0	0	0	0	0	48	0	0	48
DB2 7.1.0 5675-D B2	24	24	0	24	24	0	24	24	0	144
DB2 8.1.0 5625-D B2	24	24	24	0	0	0	0	24	0	96
DB2 9.1.0 5635-D B2	24	24	0	0	0	0	0	24	0	72
IMS/ESA V0R1 5635-A01	5	0	0	0	0	0	0	0	0	5
IMS/ESA V8R1 5655-C56	18	13	0	0	0	0	0	0	0	31
IMS/ESA V9R0 5655-J38	16	0	0	0	0	0	0	0	0	16
IMS/ESA V9R1 5655-J38	52	4	0	0	0	0	0	0	0	56
MQM MVS/ESA V6 5655-L82	0	235	0	0	0	0	0	0	0	235
Total	252	421	54	72	72	24	120	144	37	1196

Figure 2-229. Sample BATMUG Frequency Reports

2.10.3.2.2 High-Low Hourly Usage Report

To help you better understand the standard output produced by this query, definitions and calculations involved in producing the displayed variables follow:

LONG-TERM AVERAGE OF CPU SERVICE IN MSUs: CPUAVLAC is the average number of service units consumed by the partition, during the four-hour rolling periods covered by the measurement interval.

PRODUCT MSU RATE: Contains the product's hourly MSU consumption rate.

$$\text{PRDMSUR8} = \frac{\text{MUGCPUTM} * \text{SUCONFA} * 3600}{\text{DURATION} * 1000000}$$

MUGCPUTM is the sum of TCB and SRB CPU time, in hundredths of a second, as recorded by the Measured Usage License Charging (MULC) subsystem.

SUCONFA is the physical CPU adjustment factor used to convert CPU times to SUs.

The CPU adjustment factor is calculated by multiplying the CPU Capability Adj. Factor for N CPUs (CPCMAFxx), the CPU Capability (CPCCPUCB), and the Configured CPU Count (CPCCFGCT). All are from the SMF type 89 (Usage Data) record.

DURATION is equal to 3600, which is the number of seconds in an hour.

High-Low Hourly Usage Report Four Highest vs. Four Lowest System Usage Hours						
----- PRODUCT=CICS V2 R2.0 5697-E93 Total System Usage In MSUs=281.75 -----						
PRODUCT	Day of Month	Hour of Day	System Identifier	Long-Term Average of CPU Service (MSUs)	Product MSU Rate	
CICS V2 R2.0 5697-E93	4	12	CA11	107.00	0.0064	
CICS V2 R2.0 5697-E93	4	12	XAL1	44.00	3.5287	
CICS V2 R2.0 5697-E93	4	12	XAT1	130.75	77.4821	
CICS V2 R2.0 5697-E93	4	12	XE21	.	0.0064	
-----				-----	-----	
TOTAL				281.75	81.0236	
----- PRODUCT=CICS V2 R2.0 5697-E93 Total System Usage In MSUs=278.75 -----						
PRODUCT	Day of Month	Hour of Day	System Identifier	Long-Term Average of CPU Service (MSUs)	Product MSU Rate	
CICS V2 R2.0 5697-E93	4	11	CA11	98.50	0.0855	
CICS V2 R2.0 5697-E93	4	11	XAL1	46.75	3.3279	
CICS V2 R2.0 5697-E93	4	11	XAT1	133.50	85.8176	
CICS V2 R2.0 5697-E93	4	11	XE21	.	0.0855	
-----				-----	-----	
TOTAL				278.75	89.3166	

Figure 2-230. Sample High-Low Hourly Usage Report

2.10.3.2.3 All-Hour Hourly Usage Report

To help you better understand the standard output produced by this query, definitions and calculations involved in producing the displayed variables follow:

LONG-TERM AVERAGE OF CPU SERVICE IN MSUs: CPUAVLAC is the average number of service units consumed by the partition, during the four-hour rolling periods covered by the measurement interval.

PRODUCT MSU RATE: Contains the product's hourly MSU consumption rate.

$$\text{PRDMSUR8} = \frac{\text{MUGCPUTM} * \text{SUCONFA} * 3600}{(\text{DURATION}) * 1000000}$$

MUGCPUTM is the sum of TCB and SRB CPU time, in hundredths of a second, as recorded by the Measured Usage License Charging (MULC) subsystem.

SUCONFA is the physical CPU adjustment factor used to convert CPU times to SUs.

The CPU adjustment factor is calculated by multiplying the CPU Capability Adj. Factor for N CPUs (CPCMAFxx), the CPU Capability (CPCCPUCB), and the Configured CPU Count (CPCCFGCT). All are from the SMF type 89 (Usage Data) record.

DURATION is equal to 3600, which is the number of seconds in an hour, providing an hourly MSU consumption rate.

All-Hour Hourly Usage Report					
PRODUCT	Day of Month	Hour of Day	System Identifier	Long-Term Average of CPU Service (MSUs)	Product MSU Rate
z/OS 01.08.00 5645-001	1	12	CA11	25.50	15.535
			XAD1	71.00	6.398
			XA55	18.00	1.913
			XE09	.	0.051
			XE21	.	0.239
-----	-----	-----		-----	-----
z/OS 01.08.00 5645-001	1	12		114.50	24.137
z/OS 01.08.00 5645-001	1	13	CA11	52.75	8.945
			XAD1	69.25	7.251
			XAL1	38.00	0.254
			XA55	17.50	1.898
			XE09	.	0.107
			XE21	.	0.255
-----	-----	-----		-----	-----
z/OS 01.08.00 5645-001	1	13		177.50	18.710
z/OS 01.08.00 5645-001	1			292.00	42.848
z/OS 01.08.00 5645-001	2	12	CA11	46.75	9.195
			XAD1	52.00	2.744
			XAL1	29.50	1.081
			XAT1	81.75	64.406
			XA55	12.00	1.835
			XE09	.	0.188
			XE21	.	0.152
-----	-----	-----		-----	-----
z/OS 01.08.00 5645-001	2	12		222.00	79.601
z/OS 01.08.00 5645-001	2	13	CA11	50.00	9.966
			XAD1	45.75	2.741
			XAL1	27.75	1.008
			XAT1	88.75	29.264
			XA55	12.00	1.855
			XE09	.	0.208
			XE21	.	0.150
-----	-----	-----		-----	-----
z/OS 01.08.00 5645-001	2	13		224.25	45.192
z/OS 01.08.00 5645-001	2			446.25	124.793

Figure 2-241. Sample All-Hour Hourly Usage Report

2.10.4 Subcapacity Pricing Concepts

Subcapacity pricing is IBM's methodology for charging for certain software based on the utilization capacity of the LPARs in which the products run, rather than on the MSU rating of the machine.

Workload License Charges (WLC) is an example of subcapacity pricing. This pricing method offers LPAR-based pricing for subcapacity eligible software products based on the highest four-hour rolling average utilization of the LPARs where the product ran. WLC may also be implemented in full capacity mode. Full capacity mode will mean pricing is based on the MSU rating of the machine in which the product executes.

Subcapacity pricing is based on the concept of the highest observed rolling four-hour average utilization of the LPAR or LPARs where a product runs. The simultaneous combined utilization of these LPARs is determined for each hour in the reporting period and the highest observed combined utilization is used as the basis for pricing the product.

There are many advantages and disadvantages to moving to a subcapacity pricing method, so careful planning and analysis is important. You will need to understand the prerequisites for subcapacity pricing, create a software inventory, create a capacity and growth plan, and determine your current capacity and usage for each LPAR.

See the z/OS Planning for Subcapacity Pricing guide for more information on implementing subcapacity pricing for IBM products.

2.10.5 Reporting Requirements

Review the applicable sections to determine the execution requirements for the individual Product Measured Usage reports.

- 1 - LPAR Product Consumption Report
- 2 - Combined Hourly System Usage Reports

2.10.5.1 LPAR Product Consumption Report

The following are requirements that you must satisfy to execute this report.

- o CA MICS Batch and Operations Analyzer product change SMF6820 first delivered with r12 SP1 (PSP0710).

This product change introduces two essential elements: MUGDUR and MUGMXCPU.

- o Measured Usage Global (BATMUG) File activated, with the following elements populated:

- MUGPROD - Product Name
- MUGVER - Product Version
- MUGPRDID - Product ID
- MUGQUAL - Product Qualifier
- MUOWNER - Owner/Vendor Name
- MUGDUR - Product Usage Data Duration
- MUGCPUTM - Total Product CPU Time
- MUGMXCPU - Max Product CPU Time
- CPCCFGCT - Configured CPU Count
- CPCMAF0n - CPU Capability Adj. Factor for n CPUs

Only one level is required for the CA MICS SMF product files (that is, DETAIL, DAYS, WEEKS, or MONTHS), depending on the timespan you specify during inquiry definition.

The number of cycles that are maintained at your installation for the CA MICS files determines the time frame for which you can produce Product Measured Usage reports.

2.10.5.2 Combined Hourly System Usage Reports

The following are requirements that you must satisfy to execute this report.

- o The Measured Usage Global ns Analyzer product change SMF6820 first delivered with r12 SP1 (PSP0710).

This product change introduces two essential elements: MUGDUR and MUGMXCPU.

- o The CA MICS Batch and Operations Analyzer (SMF) and the CA MICS Hardware and SCP Analyzer (RMF) DIAs have to be installed in the same unit.
- o Measured Usage Global (BATMUG) File activated, with the following elements populated:

- MUGPROD - Product Name
- MUGVER - Product Version
- MUGPRDID - Product ID
- MUGQUAL - Product Qualifier
- MUOWNER - Product Owner/Vendor Name
- MUGDUR - Product Usage Data Duration
- MUGCPUTM - Total Product CPU Time
- MUGMXCPU - Max Product CPU Time
- CPCCFGCT - Configured CPU Count
- CPCMAF0n - CPU Capability Adj. Factor for n CPUs

Only the DAYS level is required for the CA MICS SMF product files, because the query can only be executed at the DAYS by Hour timespan.

The number of cycles that are maintained at your installation for the CA MICS files determines the time frame for which you can produce Product Measured Usage reports.

- o Hardware Utilization (HARCPU) File activated, with the following elements populated:

- CPUAVLAC - Long-Term Average of CPU Service (MSUs)
- CPUMXLAC - Maximum Long-Term Average of CPU MSUs
- CPUTOMSU - Actual Number of Consumed MSUs

Only the DAYS level is required for the CA MICS RMF product file, because the query can only be executed at the DAYS by Hour timespan.

The number of cycles that are maintained at your installation for the CA MICS files determines the time frame for which you can produce the Combined Hourly System Usage reports. When executing any of these reports be sure that the BATMUG file and the HARCPU file have the same number of cycles available for reporting.

Chapter 3: PERFORMANCE ANALYSIS TOOLS

CA MICS Performance Manager Option uses CA's U.S. patented Neugents(R) technology, simulation algorithms, and other methodologies to assist you in exploring measurement data.

This chapter discusses the following Performance Analysis Tools:

- 1 - Batch Initiator Simulation
- 2 - Relative Importance
- 3 - Data Clustering Analysis

This section contains the following topics:

[3.1 Batch Initiator Simulation](#) (see page 727)

[3.2 Relative Importance](#) (see page 780)

[3.3 Data Clustering Analysis](#) (see page 798)

3.1 Batch Initiator Simulation

The Batch Initiator Simulation inquiry consists of a batch initiator model and associated reports. You can use this inquiry effectively in combination with the Data Clustering inquiry discussed in Section 3.3, to analyze batch initiators.

The following sections discuss the Batch Initiator Simulation inquiry:

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output
- 4 - Analytic Technique Tutorial
- 5 - Case Study
- 6 - Component Operation
- 7 - Requirements

3.1.1 Functional Description

The Batch Initiator Simulation inquiry provided by the CA MICS Performance Manager Option allows you to evaluate the effect of modifying the job class or initiator structure for your workload. In particular, the inquiry can be applied to the data in the DETAIL timespan of the CA MICS BATJOB and BAT_JS Files. The inquiry evaluates the initiator utilization and the input queue time of the selected jobs under the proposed initiator structure.

NOTE: The Batch Initiator Simulation inquiry is not appropriate for JES3 workloads. You should use it only for JES2 workloads.

The inquiry is implemented as a discrete simulation. This offers several advantages over a queuing model approach to the problem. Principal among these is the fact that the inquiry need not make any assumption about arrival rates (for example, arrival rates need not be uniform). Other simplifying assumptions made for the purpose of the inquiry are discussed in Section 3.1.2.

The Batch Initiator Simulation inquiry produces the following decision support reports:

- o Initiator Busy Report - Shows the number of jobs processed by each initiator and the percent of time that the initiator was occupied. If the initiator serves multiple classes, figures are provided by class and for the entire initiator.
- o Job Queue Size Report - Provides queue statistics for each reporting interval of the simulation. The maximum queue length, queue length at the end of the interval, and the number of arrivals is provided for each class.

- o Job Input Queue Time Report - Consists of a plot for each class. Each plot presents a comparison of the actual and simulated queue times for jobs in that class.
- o Job Class Throughput Report - Presents a comparison of actual and simulated throughput for reporting periods. Average queue times are also provided, and all figures are broken down by job class.

These reports are described in greater detail in Section 3.1.3.

3.1.2 Usage Guidelines

The Batch Initiator Simulation inquiry operates under the following assumptions.

- o Job execution times and queue times are rounded to the next multiple of the value specified in the clock update exit (documented in Section 3.1.6.2.4). If you do not specify a value using the exit, the inquiry defaults to the largest value that is less than or equal to the 20th percentile of the incoming data.

This means that simulated data is pessimistic for initiators that serve a high percentage of very short jobs.

- o The inquiry assumes that the runtime of each job under the new initiator structure will be approximately equal to the execution time in the original environment. Because job execution times are rounded up as described above, the inquiry allows for small increases in runtime.
- o The inquiry assumes that the number of initiators in the simulated environment is less than or equal to the number of initiators in the measured environment. This assumption minimizes the potential of significant changes in the characterization of the job mix.

If you expect the mix of jobs to change, use the Specify Job Input exit documented in Section 3.1.6.2.2 to create new job classes. This allows you to determine the impact of the new job mix.

- o The initiator structure proposed will not result in resource demands that exceed the available number of physical resources like mountable tapes or disks.
- o When reviewing the reports produced by the simulation, remember that the jobs processed by the simulation inquiry are not subject to external dependencies or operator holds like those experienced by production jobs.
- o The simulation inquiry assumes that two different jobs with the same job name will not be running in the system at the same time. If the simulation inquiry attempts to take a job from the job queue and assign it to an initiator and finds that a job with the same name is already assigned to an initiator, it assumes the second job to be a duplicate and temporarily skips it. When the first job completes its execution, the second job becomes eligible for selection. This mechanism is the same one used by JES2 for managing duplicate job names.

3.1.3 Standard Output

The decision support reports produced by the Batch Initiator Simulation inquiry are discussed in the following sections:

- 1 - Common Report Elements
- 2 - Initiator Busy Report
- 3 - Job Queue Size Report
- 4 - Job Input Queue Time Report
- 5 - Job Class Throughput Report

3.1.3.1 Common Report Elements

Each report has several lines of title information that give the name of the report, the installation name, the run date, the date(s) being reported, and the start and end times of the period covered by the report. The following elements are in the standard title part of each report:

Installation Name: The installation name as defined in the SITE member of the prefix.MICS.PARMS library. The name is centered in the first line of the report.

SYSTEM: The system identifier you filled in for the Select SYSID field on the Batch Initiator Simulation inquiry panel.

Run Date: The report execution date presented in SAS WORDDATE. format.

Start Date/Time: The earliest date and time included in the report. You specified these values on the Batch Initiator Simulation inquiry panel. The Simulate date field and the first blank of the Simulate time field contain these values.

End Date/Time: The latest time included in the report. You specified this value on the last blank of the Simulate time field.

3.1.3.2 Initiator Busy Report

The Initiator Busy Report, illustrated in Figure 3-1, shows the number of jobs for each class that executed using each initiator during the report period. In addition, for each job class, the report shows the percentage of the simulation period the initiator devoted to that class.

NOTE: This report presents data produced by the simulation. The figures shown are the output of the simulation software, not actual numbers of real jobs and percent busy times.

```

                                INITIATOR BUSY REPORT
SYSTEM = TEST                      RUN DATE = NOV 19, 1987
                                MODEL PARAMETERS
START DATE/TIME                  END DATE/TIME
18NOV87:07:00                  18NOV87:12:00
                                PRIORITY AGING
                                PRILOW PRIHIGH PRIRATE
                                5      11      24

    INIT  1 CLASS: D
    CLASS  JOBS STARTED  % BUSY
    -----
    D      30           62.1
    INIT  2 CLASS: D
    CLASS  JOBS STARTED  % BUSY
    -----
    D      13           49.5
    INIT  3 CLASS: D
    CLASS  JOBS STARTED  % BUSY
    -----
    D       8           73.1
    INIT  4 CLASS: D
    CLASS  JOBS STARTED  % BUSY
    -----
    D       5           58.1
    INIT  5 CLASS: D
    CLASS  JOBS STARTED  % BUSY
    -----
    D       2           14.3
    .
    .
    .
```

Figure 3-1. Sample Initiator Busy Report

REPORT OPTIONS

The fields available on the Batch Initiator Simulation Inquiry panel, which is shown in Figure 3-11, determine the content of the information in the report. The panel and field descriptions is documented in Section 3.1.6.

DATA ELEMENTS

The elements that create the report heading are documented in Section 3.1.3.1, Common Report Elements. Other report data elements in the Initiator Busy Report are:

Priority Aging:	The PRILOW, PRIHIGH, and PRIRATE elements correspond to the Aging (low/high/rate) field on the Batch Initiator Simulation inquiry panel.
Class:	Corresponds to the Report class field on the Batch Initiator Simulation inquiry panel.
Jobs Started:	The number of jobs of the listed class that the simulation started on the initiator during the reporting period.
% Busy:	The percentage of the reporting period during which this initiator was occupied by a job of this class.

NOTE: The count of jobs is made when the job is started; the percent busy count is accumulated while the job occupies the initiator. Therefore, it is possible to have a report entry showing zero jobs but a non-zero percent busy. This situation is likely to occur if the report covers a subset of the time being simulated.

3.1.3.3 Job Queue Size Report

The Job Queue Size Report describes, for each queue and for each reporting interval, the queue of jobs waiting for batch execution for the simulation period. A sample report is shown in Figure 3-2.

NOTE: This report presents data produced by the simulation. The figures shown are the output of the simulation inquiry, not the description of actual queues.

JOB QUEUE SIZE REPORT												
SYSTEM = TEST						RUN DATE = NOV 19, 1987						
MODEL PARAMETERS												
START DATE/TIME						END DATE/TIME						
18NOV87:07:00						18NOV87:12:00						

	CLASS A			CLASS B			CLASS C			CLASS D		
TIME	MAX	END	NO	MAX	END	NO	MAX	END	NO	MAX	END	NO
	QUE	QUE	ARR	QUE	QUE	ARR	QUE	QUE	ARR	QUE	QUE	ARR
-----+-----+-----+-----+-----+-----+-----												
7:00	1	0	1	.	.	.	.	.	.	1	0	2
7:05	0	0	0	.	.	0	1	0	2	1	0	1
7:10	0	0	0	.	.	0	2	0	2	1	0	1
7:15	0	0	0	2	1	5	1	0	1	1	0	3
7:20	0	0	0	11	1	13	3	1	3	0	0	0
7:25	1	1	1	1	1	0	2	1	1	1	0	1
7:30	1	1	0	2	1	1	2	1	1	1	0	1
7:35	1	1	0	5	4	4	1	1	0	14	14	14
7:40	1	1	0	5	4	2	2	1	2	15	14	2
7:45	2	1	1	4	4	0	1	1	0	14	14	0
.	.	.	.	.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.	.	.	.	.	.
11:35	22	20	2	88	82	6	33	32	1	29	27	3
11:40	23	22	5	86	86	4	33	30	2	27	27	0
11:45	24	23	3	90	86	5	31	29	2	27	27	0
11:50	26	25	8	91	89	10	31	30	3	28	28	1
11:55	28	26	8	90	89	3	31	30	2	29	28	2
12:00	27	26	1	90	89	1	31	31	1	28	28	0

Figure 3-2. Sample Job Queue Size Report

REPORT OPTIONS

The fields available on the Batch Initiator Simulation inquiry panel, shown in Figure 3-11, determine the content of the information in the report. The Inquiry panel and field descriptions is documented in Section 3.1.6.

DATA ELEMENTS

The elements that create the report heading are documented in Section 3.1.3.1, Common Report Elements. Other data elements in the Job Queue Size Report are:

Time: The report interval covered by this report. The length of this interval is determined by the Report intvl. field on the Batch Initiator Simulation inquiry panel.

Class: Corresponds to the Report class field on the Batch Initiator Simulation inquiry panel.

The following three measures are provided for each reporting interval for each class. If no jobs for a particular class have arrived during the period of the simulation, these values are listed as missing (indicated by a period (.) in the report).

MAX QUE: The maximum queue length during the reporting interval.

END QUE: The queue length at the end of the reporting interval.

NO ARR: The number of arrivals of jobs of that class during the reporting interval.

3.1.3.4 Job Input Queue Time Report

The Job Input Queue Time Report, shown in Figure 3-3, is a set of plots, one for each class, showing both the actual and simulated queue times. The queue times are shown as cumulative distributions; the data points for the actual and simulated data are overlaid on the same plot.




Figure 3-3. Sample Job Input Queue Time Report

REPORT OPTIONS

The fields available on the Batch Initiator Simulation inquiry panel, which is shown Figure 3-11, determine the content of the information in the report. The Inquiry panel and field descriptions is documented in Section 3.1.6.

DATA ELEMENTS

The elements that create the report heading are documented in Section 3.1.3.1, Common Report Elements.

The report plots the range of input queue times for the class against the percentage of jobs that arrived for that class. This means that a point on the plot maps a given input queue time value to the percent of jobs that had an input queue time of that value or less. For example, in the report shown in Figure 3-3, 100 percent of the simulated queue times were 2:05:00 or less.

3.1.3.5 Job Class Throughput Report

The Job Class Throughput Report, shown in Figure 3-4, tabulates the number of jobs completed for each class during the period reported. The report also provides an average queue time for each class.

```

JOB CLASS THROUGHPUT REPORT
SYSTEM = XA01                      RUN DATE = SEP 17, 1987
MODEL PARAMETERS
START DATE/TIME                    END DATE/TIME
07SEP87:06:00                      07SEP87:10:00
PRIORITY AGING NOT SIMULATED

----- ACTUAL ----- ---- SIMULATED ----
JOBS  AVG-WAIT  CV  JOBS  AVG-WAIT  CV
A      17   0:02:20  1.6   18   0:00:02  0.7
B      24   0:00:08  0.8   24   0:00:17  4.3
C      19   0:00:08  0.6   20   0:00:02  0.8
D      39   0:02:22  3.4   28   0:11:46  1.3
E      10   0:02:37  2.3    0   0:00:00  0.0
F       1   0:07:52  0.0    1   0:00:04  0.0
V       6   0:00:44  1.7    0   0:00:00  0.0

```

Figure 3-4. Sample Job Class Throughput Report

REPORT OPTIONS

The fields available on the Batch Initiator Simulation Inquiry panel, which is shown in Figure 3-11, determine the content of the information in the report. The panel and field descriptions is documented in Section 3.1.6.

DATA ELEMENTS

The elements that create the report heading are documented in Section 3.1.3.1, Common Report Elements. Other data elements in the Job Class Throughput Report are:

Class:	Corresponds to the Report class field on the Batch Initiator Simulation inquiry panel.
Jobs:	The number of jobs completed during the reporting interval. For the actual job count, the end timestamp (ENDTS) is used to determine when a job ended.
Avg Wait:	Average input queue time for all jobs that ended during the reporting interval.
CV:	The coefficient of variation of the average input queue time for all jobs that ended during the reporting interval.

3.1.4 Analytic Technique Tutorial

An initiator is a multiclass server that selects jobs for processing by an MVS system. By specifying the number of initiators and the job classes each serves, you can limit the number of jobs given access to the system (that is, the system's multiprogramming level), provide different levels of service to various job classes, and allow the operating system to perform the majority of the job scheduling functions.

Because the behavior of an initiator structure that is comprised of many multiclass initiators is difficult to predict, the CA MICS Performance Manager Option includes an initiator simulation that allows you to determine the behavioral characteristics of alternative initiator structures at your installation.

One of the most extensive uses made of initiator structures is to provide various levels of service to job classes in testing environments. The model provided by the CA MICS Performance Manager Option defines service as the length of time between when a job arrives and when it is selected for execution (that is, the input delay).

You should not use the model to evaluate initiator structures for production jobs because the flow of production jobs is constrained by dependencies and operator-imposed holds rather than by the availability of initiators.

In developing the model, one important decision was to select the level of detail at which the initiator process would be modeled. Basically, this was a choice of whether or not to model the execution (server holding time) of jobs in detail.

Using highly complex queuing models, other investigators have concentrated on the effects of job interactions on execution time. However, in a testing environment, job execution times are often quite small in comparison to their input queuing delay.

To allow the investigation to concentrate on modeling the input queuing delay phenomenon, the model provided by the CA MICS Performance Manager Option assumes that the execution time of a job is unchanged when executed under a new initiator structure with a multiprogramming level less than or equal to the environment in which the jobs were originally executed.

Although this assumption initially seemed quite weak, investigations of average execution times of a number of resource-limited job classes revealed only small variations in execution time for a variety of multiprogramming levels. Moreover, actual experiences with the model, which are discussed in the case study in Section 3.1.5, have shown the effects to be negligible when the multiprogramming level decreases or is unchanged, provided that each class consists of jobs that have relatively small, homogeneous resource demands.

The input data for the model is selected from the DETAIL timespan of the CA MICS BATJOB and BAT_JS Files. These files contain one record for each job executed by the system. The model uses the following data elements:

- o JOB - job name
- o JOBPRTY - the priority at which the job was submitted
- o JOBCLASS - job class
- o JOBINQTM - job input queue time, T1
- o JOBENQTM - job enqueue time, T2
- o JOBALCTM - job allocation time, T3
- o JOBEXCTM - job execution time, T4
- o ENDTS - execution end time stamp, T5

When the job records for a system are sorted in ascending order by reader time, they provide a stream of arrivals for driving a simulation model. A job's actual queuing delay is T1, server holding time is measured as T2+T3+T4, and actual throughput is estimated using T5. As previously stated, the simulation model assumes that a job's server holding time remains unchanged, and the model determines the queuing delay under the new initiator structure.

Based on the simulation clock step value (see Section 3.1.6.2.9), job execution times are rounded. This makes the results of the model somewhat pessimistic and reduces the execution time of the model. Moreover, rounding the execution times upward elongates the execution time of each job by a few percent, allowing for small changes in execution time that may result from new job mixes.

There are a number of influences on initiator and job activity that are not recorded in the SMF log file and thus cannot be accurately reproduced by the model. These influences are introduced by operator and user interactions with the system. For instance, the operator or user may:

- o Modify the priority of a job to change its position in the input queue.
- o Put a job into a hold status to prevent or delay its execution.
- o Modify the initiator structure.

Generally, these operator and user interactions are minimal in a well-managed system. For any given system, you can determine the usual extent of these activities by investigating the console log where operator actions are recorded.

3.1.5 Case Study

This section presents an example using the Batch Initiator Simulation inquiry. The example is a validation run of the model against an actual initiator structure. It is presented in the following sections:

- 1 - Problem Description
- 2 - Parameter Specification
- 3 - Report Interpretation

3.1.5.1 Problem Description

Our sample site is running a validation run of a model against its actual initiator structure. You should always use a validation run as the first step in applying a model to a system, because the output of the validation run provides insight into how well the model simulates the type of arrival pattern and mix of priorities that occur for each job class in the workload. More specifically, the run helps you detect features of your workload (if any) that the model handles poorly.

You should not underestimate the importance of this step. Using the model without this step is similar to trying to use a map to get to point B without first determining your current position.

One of the most important decisions in model validation is the selection of the time period that you want to simulate. The objective is to select a period that represents your workload satisfactorily. Represent does not necessarily mean that some sort of average picture of the workload is the desired input. Rather, the goal is to represent the workload with both stable periods and periods when more spectacular events occur. For example, if there are peaks in the arrival rate for certain classes at certain times, these peaks should be represented in your model input.

There is no set method for developing a good set of input for the validation run. However, the task should include the steps described below. Additional steps may be required, depending on how the system being studied is used.

Characteristics of a job class that affect the model's ability to model the job class are the arrival pattern of jobs, the length of the job execution, and the priorities of the jobs. The steps below are intended to help you gain an understanding of these items for the system.

- o You must acquire an understanding of how job classes and priorities are defined and used on the system being studied. For example:
 - Are certain classes used only for production or test?
 - Are some classes used only at certain times?
 - Are some classes used only by some users?
 - What resource constraints (if any) are associated with each class?

- How are job priorities used to control the workload?
- o Consider cyclic patterns in the workload. This should include daily, weekly, monthly, and even yearly patterns if they are strong. Strength here is measured in terms of the effect on job arrival patterns and job execution times. A clustering of the workload for the system would be useful in performing this task. In addition, such areas as operating procedures for production jobs and procedures controlling test jobs can be useful in identifying important features of the system.

Once you have an understanding of the characteristics that you believe best represent the batch environment that you are modeling, choose data that has those characteristics. In choosing data, consider the following:

- o You may need to select data from several different days in order to get a good representation of the workload.
- o At a minimum, you should run the model against data for a full day. Or, if you use different initiator structures at different times of the day, you should run the model against the period in which the structure you are studying is being used.
- o If possible, choose a starting time when the input data is selected with all initiators unoccupied and queues empty.

If this is not possible, exclude the first few hours of the simulation from the reports (except perhaps in a first validation run). If these hours are not excluded, the simulation results will be distorted by jobs that arrived before the simulation period and were not finished when the simulation period began, because these jobs are not simulated.

- o Save the data that you select so that you can use the same input to simulate proposed changes to the initiator structure. This aids in before and after type comparisons, avoids the problem of losing the data because it is aged off the CA MICS database, and allows you to run the simulation without interfering with the regular database operation. Section 3.1.6.2.6 describes the Save selected input

observations parameter of the Extended Options Menu.

3.1.5.2 Parameter Specification

Figure 3-5 shows the Batch Initiator Simulation Inquiry panel for this case study.

```

/----- Batch Initiator Simulation -----\
|Command ==>
|
|Composing CA MICS Inquiry: BAT05 - Batch Initiator Study
|
|Time-span:    DETAIL                                Cycle(s) ==> 01 - __
|Data base ==> B (SECOND)    _      _      _      _
|
|Select SYSID   ==> TEST _____
|Simulation date ==> 07JAN88 (ddmonyy) Selection hours ==> 07 -- 17    (hh)
|Simulate class ==> ABCDFW__      Simulate time   ==> 07:00 17:00 (hh:mm)
|Report class   ==> _____      Report time    ==> _____ (hh:mm)
|Aging (low/high/rate) ==> 5    11  24  Report intvl. ==> 30      (1-60 min)
|
|Initiators ==> D      D      D      D      D      D      D
|                D      D      C      C      C      C      C
|                C      C      C      C      B      B      B
|                B      B      B      B      B      B      B
|                A      A      A      A      A      A      A
|                A      A      F      F      F      W      W
|                W      W      W      W      W      W      _____
|
|Specify extended options ==> Y (Y/N)
|
\-----/

```

Figure 3-5. Case Study - Batch Initiator Simulation Inquiry Panel.

The characteristics of the simulation are as follows:

- o The sample system has an initiator structure made up of 48 initiators, each of which serves a single job class.
- o The data for the simulation comes from the CA MICS database unit with an ID of B. The data comes from the 01 cycle of the BATJOB and BAT_JS files.
- o Aging will be performed for jobs with priorities 5 to 11, and jobs are eligible for aging after one hour in the queue at a given priority.
- o Only data for the system TEST and the hours 7:00 through 17:59:59.99 is selected because the initiator structure being simulated is in use only during the prime shift. The initiators are routinely drained on this system at 7:00, so the input data is selected starting at 7:00. This ensures that the simulation starts with all initiators unoccupied, which improves the accuracy of the simulation.
- o The simulation is limited to data from January 7, 1988, between the hours of 7:00 and 17:00 for classes A, B, C, D, F, and W.
- o The finest level of data kept for reporting is at 30-minute intervals.
- o The files from which the input data is selected will be written to the file PERDDOUT.BAT05.

Because we entered Y (Yes) in the Specify extended options field on the Batch Initiator Simulation inquiry panel, the Extended Options Menu was displayed. (Section 3.1.6 provides a complete description of all available extended options.) This enabled us to enter Y (Yes) in the Save selected input observations field, which brought us to the Save selected input observations panel. Figure 3-6 shows the entries we made in the Save selected input observations panel to save our input data.

```

/----- Save selected input observations -----\
| Command ==>                                     |
|                                                    |
| Inquiry Step: BAT - Batch Initiator Simulation    |
|                                                    |

```

```
| Equivalent CA MICS file(s): BATJOB          |
| Saved SAS file name(s) ==> BAT05          |
| SAS data set name      ==> MY.SAS.DATASET |
|                                           |
|                                           |
|                                           |
|                                           |
|                                           |
|                                           |
|                                           |
|                                           |
|                                           |
|                                           |
|-----|
```

Figure 3-6. Case Study - Save Selected Input Observations Panel

3.1.5.3 Report Interpretation

When the parameters presented on the panels in Section 3.1.5.2 are executed, the reports shown in Figures 3-7, 3-8, 3-9, and 3-10 are produced. Since the purpose of this run is to validate the model, the Job Input Queue Time Report and Job Class Throughput Report are the most important because they give a comparison of the actual and simulated data. The other two reports are produced from the simulation results. Each report is discussed next.

INITIATOR BUSY REPORT

The principal use of the Initiator Busy Report shown in Figure 3-7 is to see how busy each initiator is during the simulation period. In addition, the report provides a breakdown of the workload by class for initiators that serve several classes.

Initiators which are lightly used (under 25 percent) are candidates for elimination. Figure 3-7 shows that for the initiator structure used by the sample system, none of the initiators are heavily loaded; indeed, many are very lightly loaded.

This same initiator structure is used in the application strategy, where changes in the structure are tested. The fact that so many initiators are lightly loaded is central to the changes proposed.

When studying this report, also review the breakdown of occupancy time by class for heavily loaded initiators. This breakdown can be useful in detecting situations where one class is being blocked by other classes and, therefore, is not being adequately served by the initiator structure. This is not a consideration for the sample system, because each initiator serves a single class.

```

INITIATOR BUSY REPORT
SYSTEM = TEST                      RUN DATE =  JAN 8, 1988
                                MODEL PARAMETERS
START DATE/TIME                  END DATE/TIME
07JAN88:07:00                   07JAN88:17:00

```

```

PRIORITY AGING
PRILOW PRIHIGH PRIRATE
    5      11      24

```

```

INIT  1 CLASS: D
CLASS JOBS STARTED  % BUSY
-----
D      23          77.7

```

```

INIT  2 CLASS: D
CLASS JOBS STARTED  % BUSY
-----
D      22          64.7

```

```

INIT  3 CLASS: D
CLASS JOBS STARTED  % BUSY
-----
D      26          48.4

```

```

INIT  4 CLASS: D
CLASS JOBS STARTED  % BUSY
-----
D      14          30.8

```

```

INIT  5 CLASS: D
CLASS JOBS STARTED  % BUSY
-----
D      17          19.8

```

```

INIT  6 CLASS: D
CLASS JOBS STARTED  % BUSY
-----
D       4          32.9

```

```

INIT  7 CLASS: D
CLASS JOBS STARTED  % BUSY
-----
D       6          21.0

```

```

...etc...

```

Figure 3-7. Case Study - Initiator Busy Report

JOB QUEUE SIZE REPORT

The Job Queue Size Report shown in Figure 3-8 serves as a check of the information in the Job Input Queue Time Report by showing fluctuations in the job queues over the simulation period.

The report for the sample system shows the following:

- o The start of the simulation was a period of little or no batch activity (no jobs arrived in the system). This is desirable; it means that the simulation will be more accurate.
- o Classes A, B, C, and D were active. Classes F and W had little activity (these are only used for special purposes on the system being studied).
- o No class was subject to particularly large fluctuations in workload arrival rates.
- o Several classes show spikes in arrivals (e.g., 15:00 for class B), but the spikes are generally worked off in the next few periods. Class D shows some sluggishness in this regard and should be investigated further.

The Job Input Queue Time and Job Class Throughput Reports are of particular interest because they contain both actual and simulated data. In addition, both reports contain data on queue times (one in a plot and the other as an average). The two reports are discussed separately below, followed by some comments on comparing the information in the two reports.

JOB QUEUE SIZE REPORT

SYSTEM = TEST RUN DATE = JAN 8, 1988

MODEL PARAMETERS

START DATE/TIME END DATE/TIME

07JAN88:07:00 07JAN88:17:00

PRIORITY AGING

PRILOW PRIHIGH PRIRATE

5 11 24

	CLASS A			CLASS B			CLASS C			CLASS D			CLASS F			CLASS W		
TIME	MAX	END	NO	MAX	END	NO	MAX	END	NO	MAX	END	NO	MAX	END	NO	MAX	END	NO
	QUE	QUE	ARR	QUE	QUE	ARR	QUE	QUE	ARR	QUE	QUE	ARR	QUE	QUE	ARR	QUE	QUE	ARR
7:00	2	1	2	1	0	1	3	0	6	2	1	2	.	.	.	1	0	4
7:30	1	1	0	1	0	2	2	2	2	2	2	2	.	.	0	1	0	1
8:00	2	1	3	1	0	7	3	3	1	4	2	11	.	.	0	0	0	0
8:30	2	1	5	8	0	15	10	4	21	4	2	11	.	.	0	0	0	0
9:00	4	3	9	6	5	13	8	7	7	4	2	2	1	0	1	0	0	0
9:30	4	3	3	18	18	20	8	8	2	12	11	15	1	0	1	0	0	0
10:00	11	10	19	27	26	17	17	16	10	16	15	13	0	0	0	0	0	0
10:30	12	11	11	38	37	14	21	20	10	20	19	7	0	0	0	0	0	0
11:00	16	15	24	45	38	17	25	20	8	33	31	21	0	0	0	0	0	0
11:30	20	18	16	43	38	11	24	23	6	44	43	15	0	0	0	0	0	0
12:00	21	5	39	40	13	18	36	19	30	44	32	3	1	0	1	1	0	1
12:30	7	5	10	24	21	30	22	19	7	33	32	3	1	0	1	0	0	0
13:00	7	5	26	27	23	20	21	19	10	34	32	5	0	0	0	0	0	0
13:30	9	8	16	36	35	30	23	23	7	34	33	5	0	0	0	0	0	0
14:00	17	14	19	41	38	17	25	24	8	36	34	10	0	0	0	0	0	0
14:30	25	24	28	51	50	28	36	36	15	37	35	9	0	0	0	0	0	0
15:00	42	40	35	92	89	55	48	48	20	39	38	12	0	0	0	0	0	0
15:30	42	35	21	91	86	15	55	54	12	54	54	18	1	0	2	0	0	0
16:00	36	25	20	92	91	9	63	63	12	67	65	17	0	0	0	0	0	0
16:30	26	20	8	93	82	8	63	61	7	69	67	10	0	0	0	0	0	0
17:00	21	17	1	82	82	0	61	61	0	67	67	0	0	0	0	0	0	0

Figure 3-8. Case Study - Job Queue Size Report

JOB INPUT QUEUE TIME REPORT

The Job Input Queue Time Report shown in Figure 3-9 is organized by queue time. Strong agreement of the simulated and actual distributions in the report indicates that the model accurately simulates the system being studied. Discrepancies indicate problems in the model for particular classes or for the entire system.

We can make the following comments about the plots in the example:

- o The plot shows a high tail for the actual queue time. This usually reflects some external interference such as an operator hold. This lengthens the actual queue time but is not reflected by the simulation. If the tails of the distributions diverge sooner than about the 95th percentile, you should investigate the causes and be cautious in using the simulation results.
- o The number of observations hidden are caused by two factors: points at which the two plots coincide and points at which several observations occur for one of the plots (either actual or simulated). The large number of hidden observations means that the simulation closely reflects the actual system.




Figure 3-9. Case Study - Job Input Queue Time Report

JOB CLASS THROUGHPUT REPORT

As with the Job Input Queue Time Report, strong agreement of the simulated and actual data in the Job Class Throughput Report, shown in Figure 3-10, is desirable, because it means that the simulation closely reflects the actual system.

We can make the following comments about the report:

- o Classes with jobs that are very short may not be well simulated by the model, depending on the value specified for the simulation clock step.

In our sample, the simulation clock step value is 30 seconds. As the report shows, the actual times for job classes F and W are not close to the 30-second value; they are not well simulated by the model.

- o The average queue time, while reasonable for some classes, is not in particularly close agreement overall for the simulated and actual data. This indicates that further investigation of the system under study is necessary.

An analysis of other variables in the BATJOB File showed JES priority aging taking place at a rate which did not agree with the definition used in the model.

In addition, the average is a poor descriptor of the distribution of queue times for this data. This is discussed below.

COMPARISON OF QUEUE TIME AND THROUGHPUT REPORTS

Since both the Job Input Queue Time and Job Class Throughput Reports contain information describing the input queue time, it is important to compare the data in the two reports.

In our example, the comparison is striking. The throughput report shows queue times in the range of 0 to 30 minutes. The plots in the queue time report reveal that few, if any, jobs have queue times near the average. The majority of jobs have little or no queue time at all. Those that do have a queue time have times much longer than average. In short, the distribution is quite skewed; the average is not a good descriptor of the jobs being studied.

In general, if the value for the coefficient of variation is

greater than 3 on the Job Class Throughput Report, the average input queue time may be an unreliable measure due to the high degree of variability of the data. Further investigation of the system under study may be needed prior to drawing conclusions.

```

                                JOB CLASS THROUGHPUT REPORT
SYSTEM = TEST                      RUN DATE =  JAN 8, 1988
                                MODEL PARAMETERS
START DATE/TIME                  END DATE/TIME
07JAN88:07:00                   07JAN88:17:00

                                PRIORITY AGING
                                PRILOW PRIHIGH PRIRATE
                                5      11      24

----- ACTUAL -----  ---- SIMULATED ----
JOBS  AVG WAIT  CV  JOBS  AVG WAIT  CV

A      313   0:11:05  3.1  297   0:13:59  2.9
B      306   0:29:08  2.8  265   0:25:09  2.4
C      147   0:11:01  3.3  140   0:12:54  3.5
D      123   0:16:48  3.7  120   0:11:25  2.5
F        6   0:00:07  0.2    6   0:00:00  0.0
W        1   0:00:10  0.0    1   0:00:00  0.0
```

Figure 3-10. Case Study - Job Class Throughput Report

3.1.6 Component Operation

This Component Operation Section discusses the Inquiry panels that control the operation of the Batch Initiator Simulation Inquiry.

The Batch Initiator Simulation Inquiry uses data from the Batch User Job Activity (BATJOB) and Batch User Job Suspend (BAT_JS) files as input to the simulation process. The reports generated are at the DETAIL level.

The following sections discuss component operation in detail:

- 1 - Batch Initiator Simulation Panel
- 2 - Extended Options Menu
- 3 - Executing Batch Initiator Simulations

3.1.6.1 Batch Initiator Simulation Panel

Batch Initiator Simulation is one of the Performance Manager options available through the CA MICS Workstation Facility (MWF). Chapter One of this guide explains how you use MWF to get to the Performance Manager options.

```

/----- Batch Initiator Simulation ---- Data Base Set To A \
|Command ==>|
|
|Composing CA MICS Inquiry: BAT12 - Batch Initiator Study|
|
|Time-span:    DETAIL                                Cycle(s) ==> 01 - 02|
|Data base ==> A (PRIMARY) _ _ _ _ _|
|
|Select SYSID  ==> _ _ _ _ _|
|Simulation date ==> _ _ _ _ _ (ddmonyy)  Select hours ==> _ _ _ _ _ (hh)|
|Simulate class ==> _ _ _ _ _| Simulate time ==> _ _ _ _ _ (hh:mm)|
|Report class   ==> _ _ _ _ _| Report time   ==> _ _ _ _ _ (hh:mm)|
|Aging (low/high/rate) ==> _ _ _ _ _| Report intvl. ==> 15 (1-60 min)|
|
|Initiators ==> _ _ _ _ _|
|
|_ _ _ _ _|
|_ _ _ _ _|
|_ _ _ _ _|
|_ _ _ _ _|
|_ _ _ _ _|
|_ _ _ _ _|
|
|Specify extended options ==> N (Y/N)|
|
\-----/

```

Figure 3-11. Batch Initiator Simulation Inquiry Panel

Figure 3-11 illustrates a sample Batch Initiator Simulation inquiry panel. The following is a discussion of the fields that you need to code on the panel.

- Timespan:** **DETAIL** is the only valid value for this field. You cannot change it.
- Cycle(s):** The one or more cycles from the CA MICS database that you want to use for simulation, where nn is the cycle number corresponding to the beginning (and ending) cycles for which data is to be selected.
- This parameter is required. Valid values are determined by the number of cycles of data kept in your installation's CA MICS database for the database unit specified in the database field described below. The default value is the range from cycle 1 to cycle 2.
- Generally, if the SMF data at your installation is dumped at midnight, you specify only a single cycle value. If your SMF data is not dumped at midnight, you may need to specify a cycle range to include a full day of SMF data.
- Database:** The single-character CA MICS identifier for the database to be accessed.
- This field is required. Valid values are the database unit identifiers available at your installation. Contact your CA MICS System Administrator if you do not know what database units are available, or enter a question mark (?) to view a list of your database units. The default value is the database unit containing the BATJOB and BAT-JS Files. You can select up to five databases per inquiry.
- Select
SYSID:** The logical system identifier. This field is optional. Valid values are the system identifiers denoted in the sharedprefix.MICS.PARMS(SYSID) member of the CA MICS complex. Contact your CA MICS System Administrator if you do not know what system identifiers are available at your installation.

You can simulate data from up to six SYSIDs at once. The default is to report on all systems found in the data. This field limits analysis and reporting to specific systems.

Simulation date:	The date of the data you want to use for simulation, presented in ddmonyy format. This field is required. There is no default value.
Selection hours:	The range of hours used to select data for simulation, presented in hh format. This field is optional. Valid values are 00 through 23. There is no default value.
Simulate class:	The list of job classes to be simulated. This field is optional. The default is to simulate all classes contained in the BATJOB File. If you code the Specify job input exit (one of the Extended Options Menu choices) to create a new class, do not specify a class list. The default, simulate all classes, ensures that the new class is included in the simulation. You can specify a maximum of eight classes for one simulation.
Simulate time:	The earliest and latest times, respectively, to be simulated, entered in hh:mm format. This parameter is optional. Valid values are 00:00 through 23:59. The default value is 00:00 through 23:59. The time specification for the latest time is exact. For example, if you specify 10:00, the latest time included is 10:00, not 10:59:59.99, as is the case for some CA MICS time specifications. The period chosen for simulation should start when there is little or no batch activity. The ideal situation is to start with all initiators drained and queues empty, since the model cannot simulate any job that begins execution before the simulation period.
Report class:	The list of job classes to be reported. This field is optional. There is no default value. The classes you specify here must fall within the classes specified on the Simulate class field (that is, you cannot simulate classes A, B, and C, then request a report on class X).

You cannot specify more than eight job classes.

Report time:	The earliest and latest times, respectively, to be reported on, entered in hh:mm format. This field is optional. Valid values are 00:00 through 23:59. The default value is 00:00 through 23:59. The times you specify here must fall within the times on the Simulate time parameter (that is, you cannot simulate daytime hours and then request a report of nighttime hours).
Aging (low/high/rate):	The specifications for increasing the priority of a job (i.e., aging) after it has been in the input queue for a specified period of time. Enter the JES2 priority aging parameters used at your installation. (NOTE: JES3 workload simulation is not supported by CA MICS.) This field is optional. There is no default value (that is, no aging is done). The data is entered in aa bb cc format, where: aa is the lowest priority a job can have and still be considered for aging. Jobs of lower priority will not be aged, no matter how long they have been in the queue. Valid values are numeric (0-15). bb is the highest priority a job can reach by being aged (a job must have a priority less than this value in order to be considered for aging). Valid values are numeric (0-15). cc defines how long a job must be in the queue at a particular priority before it is eligible for aging. Valid values are fractional parts of a day (a rate of 24 means that a job may be aged after 1/24th of a day (one hour) in the queue).
Report intvl.:	The number of minutes used to keep data that is being reported on. This field is optional. Valid values are integers from 1 to 60.
Initiators:	The identifier for a series of strings, each of which defines the classes served by one initiator. These fields are required. You can simulate data for up to 49 initiators for each model. There is no default value.

If you are running the simulator to validate a model, these fields should define the initiator structure used on the system that you are modeling. If you are using the simulator to try a new initiator structure, these fields should define the initiator structure that you are considering. If you want to simulate more than 49 initiators, split the initiator structure into groups of less than 50 initiators and execute separately.

NOTE: The Batch Initiator Simulation inquiry estimates the amount of job backlog depth to support based on the maximum number of classes defined for any single initiator. If your system has more than five classes in any single initiator, the "small" backlog model is selected. This may cause you to receive message PER00088 'JOBQ SIZE EXCEEDED.' Should this happen, either reduce the number of classes assigned to each initiator that has more than five classes, or subdivide the simulation into more than one run. For additional assistance, contact your CA Technical Support representative.

Specify
extended
options:

The specification for further customizing the outcome of the simulation. This field is required. Valid values are N and Y. The default value is N (No). If you select Y (Yes), the Extended Options Menu appears. This menu is described in detail in Section 3.1.6.2.

3.1.6.2 Extended Options Menu

Extended options are provided to allow you to further customize the outcome of the simulation. All options are shipped with N (No) coded as the default. To activate an option, select Y (Yes). You will be presented with panels that allow you to tailor the simulation. You can return the option to its original contents by using the R (Reset) option.

Figure 3-12 shows the Extended Options Menu.

```
/----- Extended Options Menu -----\
|Command ==>|
|          |
|Inquiry Step: BAT - Batch Initiator Simulation|
|          |
|Specify global data selection exit           ===> N (Y/N/R)|
|Specify Job Input exit                       ===> N (Y/N/R)|
|Specify Job Selection exit                   ===> N (Y/N/R)|
|Specify Clock Update exit                    ===> N (Y/N/R)|
|          |
|Specify alternate input file                  ===> N (Y/N/R)|
|Save selected input observations              ===> N (Y/N/R)|
|          |
|Specify additional element keep list         ===> N (Y/N/R)|
|Specify reports to be generated              ===> N (Y/N/R)|
|Specify simulation clock step value          ===> N (Y/N/R)|
|          |
|          |
|          |
|          |
|          |
```

Figure 3-12. Extended Options Menu Panel

The following subsections describe these options in detail:

- 1 - Specify Global Data Selection Exit
- 2 - Specify Job Input Exit
- 3 - Specify Job Selection Exit
- 4 - Specify Clock Update Exit
- 5 - Specify Alternate Input File
- 6 - Save Selected Input Observations

- 7 - Specify Additional Element Keep List
- 8 - Specify Reports To Be Generated
- 9 - Specify Simulation Clock Step Value

3.1.6.2.1 Specify Global Data Selection Exit

If you enter Y (Yes) in the Specify global data selection field on the Extended Options Menu, the SAS Statements panel for specifying global data selection, shown Figure 3-13, is displayed.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> PAGE |
|
|Inquiry Step:  BAT - Batch Initiator Simulation          |
|
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy  |
|
|Cmd  SAS Statements for:  Specifying global data selection |
| - ----- |
| _ - %MACRO SELECT; |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| _ - %MEND SELECT; |
| ***** BOTTOM OF DATA ***** |
|
\-----/

```

Figure 3-13. Specify Global Data Selection Exit Panel

The global data selection exit allows you to code the SELECT MACRO, which performs any type of selection that cannot be done with the options on the Batch Initiator Simulation inquiry panel. This MACRO is invoked before the simulation takes place (e.g., as data is being input).

For example, assume that a system in New York is processing network job entry (NJE) jobs for performance classes below 30 that are submitted on a system in Chicago (system identifier SYSB). The reader timestamps (RDRTS) for the NJE jobs from Chicago reflect the one-hour time zone difference between the two systems, resulting in an inaccurate simulation if they are not adjusted. To adjust for the differences, code the SELECT MACRO as shown in Figure 3-14.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> PAGE |
|
|Inquiry Step: BAT - Batch Initiator Simulation             |
|
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy    |
|
|Cmd  SAS Statements for: Specifying global data selection  |
|  - -----|
|  - %MACRO SELECT;|
|  -   IF PERFGRP<30;|
|  -   IF JOBINCPU='SYSB' THEN DO;|
|  -       RDRTS=RDRTS+3600;|
|  -       ENDTS=ENDTS+3600;|
|  -       TDATE=DATEPART(ENDTS);|
|  -       MONTH=MONTH(TDATE);|
|  -       DAY=DAY(TDATE);|
|  -       HOUR=HOUR(ENDTS);|
|  -       END;|
|  - %MEND SELECT;|
|***** BOTTOM OF DATA *****|
|
|
\-----/

```

Figure 3-14. SAS Statements for Specifying Global Data Selection

3.1.6.2.2 Specify Job Input Exit

```

/----- SAS Statements -----\
|Command ===>                               Scroll ===> PAGE|
|
|Inquiry Step:  BAT - Batch Initiator Simulation
|
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy
|
|Cmd  SAS Statements for:  Specifying Job Input exit
|
| - -----
|_ - %MACRO PERJIN;
|
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
|_ - %MEND PERJIN;
|
|***** BOTTOM OF DATA *****|
|
\-----

```

Chapter 3: PERFORMANCE ANALYSIS TOOLS 763

For example, assume that your installation currently allows class D jobs to use any number of tape drives, but would like to separate certain jobs that use more than six drives. The new class D would allow jobs requiring 0 to 6 drives, and a new class E would be used for jobs requiring more than six drives. You can test this new class definition by defining the PERJIN MACRO as shown in Figure 3-16.

This PERJIN MACRO selects jobs that use more than six tape drives and sets their job class to class E. Running a simulation with the new class E information allows your installation to determine the effect of adding this class on job throughput.

NOTES:

- o You should not code subsetting IFs in this MACRO to perform data selection; these should be coded in the global data selection exit described in Section 3.1.6.2.1.
- o You cannot change the reader timestamp (RDRTS) via this exit. To change the reader timestamp, use the global data selection exit. This consideration is important if you are processing NJE data for jobs that were read into a system in one time zone, but were executed on a system in a different time zone. Section 3.1.6.2.1. provides an example of using the global data selection exit for this purpose.

3.1.6.2.3 Specify Job Selection Exit

If you enter Y (Yes) in the job selection field on the Extended Options Menu, the SAS Statements panel for specifying a job selection exit, shown in Figure 3-17, is displayed.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> PAGE |
|Warning: Set SELECT=0; to skip an observation.           |
|Inquiry Step: BAT - Batch Initiator Simulation           |
|                                                         |
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy  |
|                                                         |
|Cmd  SAS Statements for: Specifying Job Selection exit   |
| - -----|
|_ - %MACRO PERJSX;|
| - -----|
| - -----|
| - -----|
| - -----|
| - -----|
| - -----|
| - -----|
| - -----|
| - -----|
| - -----|
|_ - %MEND PERJSX;|
|***** BOTTOM OF DATA *****|
|                                                         |
\-----/

```

Figure 3-17. Specify Job Selection Exit Panel

At the start of each interval, the simulation attempts to start a job for any initiators that are idle. The simulation selects the highest priority job in the first class served by the initiator with any jobs waiting.

Once this candidate job is selected, the job selection exit is invoked. It allows you to accept or reject the selected job based on criteria coded in the PERJSX MACRO. If the PERJSX MACRO sets the SELECT variable to 1, the job is accepted. If it sets the SELECT variable to 0, the job is rejected. The default sets SELECT to 1.

If a job is rejected by setting SELECT to 0, no job is selected from that queue in that interval of the simulation. If your criteria for accepting or rejecting a job is based on job priority, this will probably be of no concern, since all other jobs in the queue are of equal or lower priority. However, if you are testing some other condition, you must be sure to consider the effect of attempting selection only once for each queue served by the initiator.

Because job selection takes place before the variables of the BATJOB File are available, only the following variables are available to the PERJSX MACRO:

- o J_CLS(JP) - job class
- o J_PTY(JP) - job priority
- o J_ARP(JP) - job arrival priority
- o J_RDR(JP) - job reader time
- o J_EXT(JP) - job execution time
- o J_QUE(JP) - job queue time

The variables must be suffixed by (JP). This is the array subscript and it is required.

For example, assume that your installation is considering preventing the execution of jobs above priority 13 without the approval of the data center manager.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> PAGE |
|Warning: Set SELECT=0; to skip an observation.             |
|Inquiry Step: BAT - Batch Initiator Simulation              |
|                                                            |
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy     |
|                                                            |
|Cmd  SAS Statements for: Specifying Job Selection exit      |
| -  -----|

```

```
| _ - %MACRO PERJSX;                                     |
| _ -   IF J_PTY(JP)>=13 THEN SELECT=0;                 |
| _ -   ELSE SELECT=1;                                   |
| _ -   -----                                         |
| _ -   -----                                         |
| _ -   -----                                         |
| _ -   -----                                         |
| _ -   -----                                         |
| _ -   -----                                         |
| _ -   -----                                         |
| _ -   -----                                         |
| _ - %MEND PERJSX;                                     |
| ***** BOTTOM OF DATA *****                      |
| \-----/                                           |
```

Figure 3-18. SAS Statement for Specifying a Job Selection Exit

The PERJSX MACRO shown in Figure 3-18 accepts the job selected if its priority is less than 13.

Using this MACRO in conjunction with the clock update exit documented in Section 3.1.6.2.4 provides greater flexibility in simulations (see Figure 3-21).

The job selection exit is located at a point in the code where the system has selected a job for execution.

NOTE: You should not code subsetting IFs in this job selection exit MACRO to perform data selection. You should code them in the global data selection exit described in Section 3.1.6.2.1.

The simulation clock update exit (PERINT) is invoked each time the simulation clock is updated. You can use it to set user variables you want available during each interval of the simulation, examine the queues in detail at certain points during the simulation, dynamically change the queues, or change the initiator structure.

For example, to set the minimum priority for a job based on the current value of the simulation clock, code the PERINT MACRO as shown in Figure 3-20.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> PAGE|
|
|Inquiry Step:  BAT - Batch Initiator Simulation
|
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy
|
|Cmd  SAS Statements for:  Specifying Clock Update exit
|  - -----
|  - %MACRO PERINT;
|  -   RETAIN MINPRTY;
|  -   IF HOUR(INOW)=>0 AND HOUR(INOW)=<6 THEN MINPRTY=0;
|  -   IF HOUR(INOW)=>7 AND HOUR(INOW)=<8 THEN MINPRTY=7;
|  -   IF HOUR(INOW)=>9 AND HOUR(INOW)=<11 THEN MINPRTY=11;
|  -   IF HOUR(INOW)=>12 AND HOUR(INOW)=<13 THEN MINPRTY=9;
|  -   IF HOUR(INOW)=>14 AND HOUR(INOW)=<16 THEN MINPRTY=11;
|  -   IF HOUR(INOW)=>17 AND HOUR(INOW)=<18 THEN MINPRTY=9;
|  -   IF HOUR(INOW)=>19 AND HOUR(INOW)=<20 THEN MINPRTY=7;
|  -   IF HOUR(INOW)=>21 AND HOUR(INOW)=<23 THEN MINPRTY=11;
|  - %MEND PERINT;
|***** BOTTOM OF DATA *****|
|
\-----/

```

Figure 3-20. SAS Statements for Specifying a Clock Update Exit

The PERINT MACRO uses INOW, the variable containing the current value of the simulation clock, to set the minimum priority allowed during the time being simulated.

Using this MACRO in conjunction with the job selection exit PERJSX, documented in Section 3.1.6.2.3, allows you to determine the effects of preventing certain job classes from executing during certain time periods.

Using the output of the PERINT MACRO above, you can code the PERJSX MACRO as shown in Figure 3-21 and obtain the effect of preventing prime time execution of jobs whose priorities are less than 7.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> PAGE |
|Warning: Set SELECT=0; to skip an observation                |
|Inquiry Step: BAT - Batch Initiator Simulation              |
|                                                            |
|Line Cmds: I Insert D Delete R Repeat M Move C Copy        |
|                                                            |
|Cmd  SAS Statements for: Specifying Job Selection exit      |
| - -----                                                |
| _ - %MACRO PERJSX;                                         |
| _ - IF J_PTY(JP)>=MINPRTY THEN SELECT=1;                  |
| _ - ELSE SELECT=0;                                         |
| -----                                                  |
| -----                                                  |
| -----                                                  |
| -----                                                  |
| -----                                                  |
| -----                                                  |
| -----                                                  |
| _ - %MEND PERJSX;                                          |
|***** BOTTOM OF DATA *****                             |
|                                                            |
|-----\

```

Figure 3-21. SAS Statements for Job Selection Exit

NOTE: You should not code subsetting IFs in this clock update exit MACRO to perform data selection. You should code them in the global data selection exit described in Section 3.1.6.2.1.

Use the Specify alternate input file extended option to specify a SAS data set name and SAS file names to be used as the input source to the reporting software. When inquiries require multiple input files, you must specify all indicated CA MICS equivalent files. If you are specifying a previously saved input observation file, some files may not be needed as they are already included in the saved file. In this case, you must enter a nonexistent file name.

A common use for this extended option is in the iterative processing of an analysis program. In this scenario, an initial execution of the analysis inquiry inputs and subsets data from a CA MICS file (possibly an archive tape file when only one SAS file is required by the analysis/report), retaining the selected data through the Save selected input observations extended option. Subsequent analysis executions use this Specify alternate input file extended option to reprocess the saved data. In this way, you can economically reprocess input data against a variety of analysis parameters without revisiting the original CA MICS files.

[illegible]

The Save selected input observations extended option allows you to specify a SAS data set and file names to receive the input observations processed by the reporting software. The number of files that are created and their content are report dependent. When inquiries save data to multiple files, you must provide alternative output file names for all of the indicated files. The saved files can be reread as alternate input files in subsequent analysis inquiries.

A common use for this extended option is in the iterative processing of an analysis program. In this scenario, an initial execution of the analysis inquiry inputs and subsets data from a CA MICS file (possibly an archive tape file when only one SAS file is required by the analysis/report), retaining the selected data through this Save selected input observations extended option. Subsequent analysis executions can use the Specify alternate input file extended option to reprocess the saved data. In this way, you can economically reprocess input data against a variety of analysis parameters without revisiting the original CA MICS files.

3.1.6.2.7 Specify Additional Element Keep List

If you enter Y (Yes) in the Specify additional element keep list field on the Extended Options Menu, the Additional Element Keep List panel, shown in Figure 3-24, is displayed.

[illegible]

Figure 3-24. Additional Element Keep List Panel

The element names coded on this panel are added to the work files created by the Save selected input observations panel documented in Section 3.1.6.2.6.

The data element labels are filled in by the system; you need only fill in the element names. Valid values are the elements defined to the CA MICS database.

Normally, only variables required by the Batch Initiator Simulation inquiry are kept and subsequently output to a work file. At times, additional variables may be required for proper exit processing to occur in subsequent runs. If you include these variables as a list of variable names in this exit, they will be output to the work file and available for future processing.

3.1.6.2.8 Specify Reports To Be Generated

If you enter Y (Yes) in the Specify reports to be generated field on the Extended Options Menu, the Reports To Be Generated panel, shown in Figure 3-25, is displayed.

[illegible]

Figure 3-25. Reports To Be Generated Panel

The Batch Initiator Simulator is shipped to produce all reports. To turn off one or more reports, code an N (No) in place of the Y (Yes) on this panel.

If you catalog the inquiry, these specifications are saved. When you rerun this inquiry, only the reports with a Y are produced.

The reports are described in Section 3.1.3 of this guide.

[illegible]

778 Performance Manager Option Guide

The simulation clock step value is the multiple (in seconds) used to simulate job arrival rates and execution times.

The valid values are shown on the panel; the default value is blank, meaning that the data will be scanned to determine the largest value which is less than or equal to the 20th percentile of the incoming data. This value becomes the default.

Small values provide a more granular simulation which may improve validation. Large values, on the other hand, are less costly to run and are suitable in cases where a high degree of accuracy is not essential.

For example, setting the value to 5 causes the simulation clock step to be five seconds. This means that initiator busy status and job queue lengths, both a function of job arrival rates and execution times, are evaluated at five-second intervals. For instance, if initiator X is free at time T1 and there is a job waiting to execute, the initiator is presumed busy for the next five seconds, even though the actual job execution time may have been less. Hence, the simulated throughput of the initiator is more pessimistic than the actual initiator.

3.1.6.3 Executing Batch Initiator Simulations

When you finish specifying parameters on the Batch Initiator Simulation inquiry panel, enter the END (PF3) command. The inquiry is executed as a normal MICF inquiry, as described in Section 1.4.3.5, Executing an Analysis Inquiry.

3.1.7 Requirements

The table below shows the CA MICS files and timespan levels that are required to run the Batch Initiator Simulation inquiries to model the initiator structure at your installation.

CA MICS FILE	LEVEL				
	X	D	W	M	Y
BATJOB	x				
BAT_JS	x				

The number of cycles of the DETAIL level BATJOB File that are maintained at your installation determines the time period available for modeling. Only one cycle of the DETAIL level BAT_JS File is available.

3.2 Relative Importance

A new facility, Relative Importance, has been added to the CA MICS Performance Manager Option that will analyze the effects of a series of numeric data values (independent element) upon a client specified dependent element, and then quantify those effects. This exploratory data analysis tool can be used in selecting related data elements for future performance analysis.

The sections that follow describe the Relative Importance Facility.

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output
- 4 - Component Operation
- 5 - Case Study
- 6 - Requirements

3.2.1 Functional Description

Relative Importance (CA's U.S. patented Neugents(R) technology facility) utilizes proprietary analytical technology to understand and quantify the effect of a list of data values (independent element) against a client specified dependent element. While a similar analysis can be performed using standard statistical means such as correspondence matrix, the Relative Importance facility can not only determine linear relationships among data values but non-linear relationships as well. This is important when exploring data that may not be well understood and can result in a more accurate understanding of the relationships among the data values. As with other CA MICS Performance Manager facilities, any files to be analyzed using Relative Importance must reside in the CA MICS database.

The process itself is straightforward. Determine the dependent data element to use for analysis, and CA MICS Performance Manager will produce a default list of all quantitative independent data elements defined on the Relative Importance Analysis screen.

During execution, the Relative Importance facility will analyze the values of each selected independent element and develop a weight that quantifies the effect of the data value against the dependent element. The sum of the weights for all selected elements will be 100 and will quantify the impact of each selected independent data element against the dependent variable. Those data values with a greater weight have a more significant impact on the dependent element than those with a lesser weight, and will likely be of more interest in any future analysis.

3.2.2 Usage Guidelines

Use of Relative Importance is easy and the operation is straightforward. You can quickly execute a complete analysis with a single panel and a few key strokes. Once the analysis has completed, you can examine the output report produced, quickly determine those independent elements that have minimal effect on the dependent element values, remove those independent elements and reexecute the Relative Importance analysis again. This process can be executed as often as necessary until the best set of independent elements remains. Once an optimum set of independent elements has been determined, they can be used in future studies.

3.2.3 Standard Output

The Relative Importance Facility currently produces a single printed output: The Relative Importance Analysis Report. This report contains a set of simple descriptive statistics for the dependent elements and their associated values, as well the relative importance weight assigned to each independent individual element. A sample of this report is shown next.

RELATIVE IMPORTANCE ANALYSIS REPORT

REPORT FOR: 09/06/02

ELEMENT	MIN. VALUE	MAX. VALUE	MEAN VALUE	STD	CV	REL. IMP.
CPUADBTM	19:48:03.06	177:30:19.63	150:24:18.49	28:23:26.31	18.88	
BATCH	1:34:22.97	23:54:35.96	12:30:32.64	4:29:03.99	35.85	0.1261
CICS	4:03:04.13	39:19:11.42	33:07:25.75	6:37:24.04	20.00	0.3231
DB2	6:32:51.97	76:14:35.50	57:31:46.57	13:15:40.67	23.05	0.4012
TS0	0:02:28.15	0:54:55.72	0:17:59.79	0:07:52.44	43.75	0.1496

TOTAL RELATIVE IMPORTANCE WEIGHTS: 1.0000

The Relative Importance Analysis report contains the following information:

ELEMENT:	The name of the independent and dependent elements being analyzed.
MIN. VALUE:	The minimum value of the selected data element observed in the input data.
MAX. VALUE:	The maximum value of the selected data element observed in the input data.
MEAN VALUE:	The mean of the values of the selected data element observed in the input data.
STD:	The standard deviation of the values of the selected data element observed in the input data.
CV:	The coefficient of variance of the values of the selected data element observed in the input data. This is a standard statistic and is a

ratio of the mean and standard deviation measurements.

REL. The relative importance of this particular
IMP.: independent element as a function of the
 dependent element value. Neugents technology
 computes this value internally using a mixture
 of both linear and nonlinear techniques and
 expresses this as a ratio or weight factor,
 relative to 100.

3.2.4 Component Operation

This section describes the control parameters that are used in the execution of the Relative Importance Analysis inquiry. The control parameters let you define the CA MICS file you want to use as input, the time periods you want to analyze, the dependent elements, and the associated independent elements you want to use in the Relative Importance process.

The following sections discuss each control panel in detail:

- 1 - Relative Importance Analysis Control Panel
- 2 - Extended Options Panel
- 3 - Inquiry Processing Panel
- 4 - Data Element Derivation Panel
- 5 - Alternate Input File Panel
- 6 - Save Selected Input Observations Screen

3.2.4.1 Relative Importance Analysis Control Panel

Relative Importance Analysis is one of the Performance Manager options available through the CA MICS Workstation Facility (MWF). Chapter One of this guide explains how to use MWF to reach the Performance Manager options and subsequently the Relative Importance control panel shown in Figure 3-27.

```

/----- Relative Importance Analysis -----\
|Command ==>                               Scroll ==> CSR          |
|                                           |                      |
|Composing CA MICS Inquiry:   IOISC   - Adhoc Inquiry -          |
|CA MICS file ==> CPU - CPU ACTIVITY FILE           Time-span DETAIL |
|Data base ==> P              Allow Run-time Override of DB(s) => Y (Y/N) |
|Cycle(s) ==> 01 - 02        Allow Run-time Override of Cycles => Y (Y/N) |
|                                           |                      |
|Data Selection:                               Allow runtime override |
| Date           ==>         -             -             ==> N (Y/N) |
| Hour           ==>         -             -             ==> N (Y/N) |
| Zone           ==>                               ==> N (Y/N) |
| SYSID          ==>                               ==> N (Y/N) |
|                                           |                      |
| Specify Dependent Element => CPUTOMSU   ( Use a ? to select from list ) |
|                                           |                      |
| Specify Two or More Independent Elements Below ( Use a ? to select from list ) |
| > CPUTTT        >CPUTTT        >CPUTTT        >CPUTTT        >CPUTTT |
| >               >               >               >               > |
|                                           |                      |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 3-27. Relative Importance Analysis

Some of the parameters on this panel are required, while others are optional. Each parameter is discussed next.

CA MICS file:	The CA MICS database file containing the historical observations for the Analysis file used in the Relative Importance Analysis application. If you are uncertain of the files or file names available in your CA MICS database, type a question mark (?) in the field and a selection list will be displayed.
Database:	The single-character database identifier for this file at your site.
Cycle(s)	The number of cycles to be specified. For example, we specified two cycles to analyze two days worth of data.
Data selection:	The selection criteria for Date, Hour, Zone, and SYSID to subset the data you are examining. You must enter date values in the ddmonyy format.
Specify Dependent Element	The name of the dependent element used in the Relative Importance Analysis application. If you are uncertain of the elements available in your CA MICS file, type a question mark (?) in the field and a selection list will be displayed.
Specify Independent Elements	The independent elements that you select for the analysis. If you are uncertain of the elements available in your CA MICS file, type a question mark (?) in the field and a selection list will be displayed.
Specify extended options:	The specification for selecting extended options. Select this option if you wish to specify additional parameters.

3.2.4.2 Extended Options Panel

After the "tapeload" libraries have been copied from the distribution tape, execute this step to load the remaining CA MICS complex libraries. Edit sharedprefix.MICS.TAPELOAD.CNTL and select the MICSLD1 member. Review and complete the symbolic parameters on the EXEC statement. DO NOT change the parameters on the PROC statement.

The EXEC statement is:

```
//LOAD EXEC LOAD
//  RSIZE='6M',           <-- REGION FOR ALL STEPS
//  DSP=NEW,              <-- NEW IF COMPLEX NOT EXIST
//                          <-- OLD IF COMPLEX EXISTS
// *
//  TUNIT='_____',        <-- REQUIRED --<<
//  TEXPDT=98000,         <-- IS TAPE EXPDT OK?
//  TVOL=_____,           <-- REQUIRED --<<
//  LIBUNIT=,             <-- LIB DASD UNIT REQUIRED
//  LIBVOL=,              <-- DO YOU NEED LIB DASD VOL?
//  LIBSTORC=,            <-- OPTIONAL LIB STORCLAS=
//  LIBDATAC=,            <-- OPTIONAL LIB DATACLAS=
//  LIBMGMTC=,            <-- OPTIONAL LIB MGMTCLAS=
//  LIBDSNT=,             <-- OPT, PUT LIBRARY FOR PDSE
//  WKUNIT=SYSDA,         <-- IS WORK DASD UNIT OK?
//  BLOCK=6160,           <-- LIBRARY BLKSIZE?
//  VBLOCK=6160,          <-- BLKSIZE FOR VB LIBRARIES?
//  SPREFIX='_____',      <-- REQUIRED --<<
//  SMICS='MICS.',         <-- SHARED MICSLEVEL?
//  LOADER=IEWLDRGO       <-- YOUR SITE'S LOADER NAME
```

where:

RSIZE is the amount of region storage for steps in this job.

DSP is the disposition of the libraries being loaded: NEW if a new complex is being loaded; OLD if the complex already exists.

TUNIT, TEXPDT, TVOL are the unit, expiration date, and volume serial number of the CA MICS distribution tape.

LIBUNIT, LIBVOL, LIBSTORC, LIBDATAC, LIBMGMTC, BLOCK, VBLOCK are the unit, volume serial number, SMS storage class (STORCLAS=), SMS data class (DATACLAS=), SMS management class (MGMTCLAS=), and block size to be used for the production libraries. Note that the

value for BLOCK must be greater than or equal to the block size for the ISPF panel and table libraries (ISPPLIB and ISPTLIB). In addition, VBLOCK is used for the documentation text libraries, which are RECFM=VB format.

LIBDSNT can be either blank or LIBRARY. Specify LIBRARY to allocate new data sets as partitioned data set extended (PDSE) data sets. Any library ending with .LOAD will be allocated as a PDS.

WKUNIT is the unit for the SYSUT3 and SYSUT4 work data sets needed by IEBCOPY.

SPREFIX is the sharedprefix for the database complex. Note that this is specified with no trailing period (for example, DB1.SP rather than DB1.SP.).

SMICS is the "MICS" level of the standard CA MICS data set names. If you do not want "MICS" as part of your data set names, specify "SMICS=" to set it to null.

LOADER is the program name of the IBM loader at your site. This is typically IEWLDRGO, HEWLDRGO, or LOADER. If you are not sure, contact your system programmer.

After editing the JCL, submit the job and exit
sharedprefix.MICS.TAPELOAD.CNTL.

```
*****
*
* NOTE: YOU MUST EXIT TAPELOAD.CNTL AFTER YOU SUBMIT THE
*      JOB.
*
*****
```

Note that the PARMCOPY step may end with a condition code of 4. This is a result of members being selected for a product for which your site is not licensed. This is an acceptable condition for the PARMCOPY step only. All other steps must end with a condition code of 0.

3.2.4.3 Inquiry Processing Panel

Figure 3-29 illustrates the Inquiry Processing panel.

```

/----- Inquiry Processing -----\
|Command ==>|
|
|   Active MICF Inquiry: RICPU - Relative Importance
|           Author:      Your Name
|
|   Specify Inquiry processing ==> EXECUTE (EXECUTE/MODIFY/SAVE/CANCEL)
|   EXECUTE - Execute the inquiry.
|   MODIFY  - Redisplay the inquiry for modification.
|   SAVE    - Exit from inquiry composition and save the inquiry.
|   CANCEL  - Exit from inquiry composition without saving the inquiry.
|
|Instructions:
|
|   Specify desired option and press ENTER key.
|
\-----/

```

Figure 3-29. Inquiry Processing Panel

Typing END at the Command prompt or pressing the END PF key (usually PF3) while on the Relative Importance Analysis panel will display the Inquiry Processing panel. Your available options are described in the text of the panel. Note that you can save a partially composed inquiry for later modification and execution.

3.2.4.4 Data Element Derivation Panel

If you require user-defined data elements, use the Data Element Derivation Macro Definition panel shown in Figure 3-30 to specify their derivation.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> CSR|
|
|Inquiry Step: Relative Importance
|
|Line Cmds: I Insert   D Delete   R Repeat   M Move    C Copy
|
|Cmd  SAS Statements for: Specifying data element derivation
| - -----
|_ - %MACRO DRVMAC;
|
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
| - -----
|_ - %MEND DRVMAC;
|
|***** BOTTOM OF DATA *****|
\-----\

```

Figure 3-30. Data Element Derivation Macro Definition Panel

Relative Importance provides a user-exit for defining the procedures for computing data elements that are not part of the CA MICS database. This exit is invoked by coding SAS statements in the DRVMAC macro. This macro is bracketed by the %MACRO and %MEND statements that are field protected to prevent accidental modification. The default for the DRVMAC macro is null.

You must code each data element derivation SAS statement in accordance with the syntax of the SAS Macro Language, as described in the SAS User's Guide: Basics.

3.2.4.5 Alternate Input File Panel

Figure 3-31 displays the Alternate Input File panel.

```
/----- Alternate Input File -----\
|Command ==>
|
|Inquiry Step: Relative Importance
|
|SAS file name(s) ==> _____
|
|
|SAS Data set allocation instructions:
|   Data set name    ==> _____
|   Disposition      ==> SHR  (SHR/OLD)
|
|
|
|
|\-----
```

Figure 3-31. Alternate Input File Panel

Relative Importance has the ability to read data from CA MICS files, or from non-CA MICS external SAS files. The Alternate Input File panel lets you specify non-CA MICS SAS files as input data for a Relative Importance analysis. You can specify the SAS data set name and up to five SAS files to be used as input to the analysis. Note that the SAS data must already be created.

3.2.4.6 Save Selected Input Observations Screen

Figure 3-32 illustrates the Save Selected Input Observations screen.

[illegible]

Figure 3-32. Save Selected Input Observations Screen

Relative Importance has as the ability to save the observations that your analysis inquiry selected to an external SAS file for further processing by independent SAS programs, or for future processing by CA MICS Performance Manager. The Save Selected Input Observations screen allows you to specify that the selected observations are to be written to an external SAS file, and to specify the SAS data set name and the name of the SAS file to which they are to be written. Note that the SAS data set must have already been created.

3.2.5 Case Study

The following case study involves the use of Relative Importance Analysis used in preparation of a processor resource forecast.

- 1 - Processor Resource Usage
- 2 - Control Parameters
- 3 - Relative Importance Analysis Report

3.2.5.1 Processor Resource Usage

This case study shows how a you can perform a preliminary analysis of processor usage data and better understand which general workloads appear to influence the overall system usage. This may be thought of as a simple exploratory step that can help direct other more detailed analyses.

Since performance management is often done in a very individualized fashion and the actual metrics vary among sites, it is important to be able to have tools that are flexible and permit you to choose among a variety of metrics, measurements, and methods to achieve the best outcome.

DEVELOPMENT OF CONCEPTUAL STUDY

Knowledge gained from this and similar limited analyses can help you decide which metrics to use in more detailed studies. For example, you might decide that simple usage data, as depicted here, is too limited to be reliable, and that data from the WLMSECxx or related files would be required. But for quick studies and simple "back of the envelope" analyses, the method used here may be sufficient.

APPLICATION OF CONCEPTUAL STUDY

The systems being examined are SYS1 and SYS2 (these are the CA MICS SYSID values). These systems process varied workloads, which can be loosely grouped together as BATCH, TSO, STC, and USS. The systems are used for both system development and non-scheduled production, and serve a number of remotely located groups. The hardware platforms have changed over time and both LPARs reside on a z-900 based CEC executing z/OS Version 1.4 or higher.

The DETAIL HARCPU data from a typical day was chosen and the analyses performed. The analysis report was reviewed, and the decision was made to drop the "average" data elements (CPUAVx) in favor of the "maximum" values (CPUMXx), as stated above. The query was then re-executed and the results noted. The next sections explore the results of the report that is generated from the Relative Importance Analysis developed for the workload usage study:

- 1 - Control Parameters
- 2 - Relative Importance Analysis Report

3.2.5.2 Control Parameters

Using a short set of menus that are described in Section 3.2.4.1 of this guide, you see the Relative Importance Analysis panel shown in Figure 3-30.

Figure 3-30 shows the control parameters that were used to generate the Relative Importance Analysis query in this case study.

3.2.5.3 Relative Importance Analysis Report

Figure 3-31 shows the Relative Importance Analysis Report for this case study.

As this report shows, there is no significant difference between the "average" and "maximum" data element values, and in both groups, BATCH (CPUxxB) is the least relevant. The other workloads are about even in terms of affect on the CPUTOBTM value. The BATCH workload value is retained, and the "maximum" data values are used in the second execution.

RELATIVE IMPORTANCE ANALYSIS REPORT						
REPORT FOR: 10/17/03						
ELEMENT	MIN. VALUE	MAX. VALUE	MEAN VALUE	STD	CV	REL. IMP.
CPUTOBTM	0:09:53.34	0:53:38.18	0:26:23.91	0:12:40.76	48.03	
CPUAVB	3.19	21.80	6.26	2.44	39.05	0.0557
CPUAVE	7.00	13.68	10.13	2.92	28.79	0.1481
CPUAVS	190.21	279.31	236.91	40.26	16.99	0.1456
CPUAVT	2.30	60.70	25.12	16.93	67.39	0.1586
CPUMXB	5.00	25.00	9.57	3.77	39.36	0.0458
CPUMXE	7.00	15.00	10.69	3.36	31.43	0.1398
CPUMXS	195.00	281.00	238.62	40.08	16.80	0.1470
CPUMXT	3.00	63.00	27.20	17.26	63.47	0.1593
TOTAL RELATIVE IMPORTANCE WEIGHTS:						1.0000

Figure 3-31. Relative Importance Analysis Report - First Execution

For the second execution, only the "maximum" (CPUMXx) data values are used. Not surprisingly, the approximate ratios of the Relative Importance Index value do not change much between executions. In a real-life study, the actual processor times for the workloads would need to be considered as well as the simple number of users and a detailed workload characterization study would be required before using the metrics for any forecasting operation.

RELATIVE IMPORTANCE ANALYSIS REPORT

REPORT FOR: 10/17/03

ELEMENT	MIN. VALUE	MAX. VALUE	MEAN VALUE	STD	CV	REL. IMP.
CPUTOBTM	0:09:53.34	0:53:38.18	0:26:23.91	0:12:40.76	48.03	
CPUMXB	5.00	25.00	9.57	3.77	39.36	0.1116
CPUMXE	7.00	15.00	10.69	3.36	31.43	0.2493
CPUMXS	195.00	281.00	238.62	40.08	16.80	0.2922
CPUMXT	3.00	63.00	27.20	17.26	63.47	0.3469
TOTAL RELATIVE IMPORTANCE WEIGHTS:						1.0000

Figure 3-32. Relative Importance Analysis Report - Second Execution

3.2.6 Requirements

The exploitation of Neugents technology 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.

3.3 Data Clustering Analysis

Data Clustering using CA's U.S. patented Neugents(R) technology, provided in CA MICS Performance Manager, enables you to apply a clustering methodology to many types of analysis.

In the past Workload Characterization was the analysis tool provided with CA MICS Performance Manager. Workload Characterization is now the case study and Data Clustering, using Neugents technology, is the tool. The Data Clustering tool has been added to allow open ended analysis of any CA MICS file and any elements contained in that file. With this new feature you are not limited to a workload study, you can cluster any data in any CA MICS file.

The discussion that follows will refer to Workload Characterization using Data Clustering as the analysis tool.

Data Clustering using Neugents technology, provided in CA MICS Performance Manager, enables you to apply a clustering methodology to the analysis of your installation's workload. Data Clustering is an attractive approach to performance management problems because it allows the number of workload elements that need be considered in a study to be reduced from tens of thousands to only a few.

Clustering methods are a statistical extension of scatter plots to identify similarities and differences between workloads. Scatter plots are often difficult to prepare and depend heavily on visual interpretation of the data. The need for visual interpretation limits the use of scatter plots to two or perhaps three axes. Clustering overcomes these disadvantages through its ability to recognize patterns in multiple dimensions. Using the CA MICS database as an input data source for clustering simplifies and extends the application of the technique.

Note: Neugents technology is a prerequisite for using Data Clustering.

The following sections describe the analytic techniques, operational considerations, and reports used in the application.

- 1 - Functional Description
- 2 - Usage Guidelines
- 3 - Standard Output
- 4 - Analytic Technique Tutorial

- 5 - Component Operation
- 6 - Case Studies
- 7 - Requirements

3.3.1 Functional Description

During the early 1970s, Domenico Ferrari introduced the concept of workload characterization (FER72). The basic premise of the concept was that studies of computer systems could be facilitated by the identification of a number of "resource patterns" that could be used to statistically represent a system's workload. Since capacity planning and performance management studies often encompass tens of thousands of jobs or hundreds of thousands of transactions, the ability to represent the workload by a limited number of resource patterns can significantly simplify a study. (Typical studies result in the identification of from 10 to 30 patterns.)

Ferrari identified these resource patterns, which he called clusters, by visual inspection of two-dimensional scatter plots. For example, Ferrari reported studies where he examined scatter plots of print lines and CPU seconds representing the resource usage patterns of student jobs. Unfortunately, the dependence on visual inspection limited the technique to two or perhaps three dimensions, since graphic representations become extremely complex for higher dimensions. Also, the technique depended upon the visual skill of the interpreter, since some analysts could "see" patterns better than others. These problems limited the applications of what was otherwise a valid concept.

In 1976, however, two papers were published that detailed the application of statistical pattern recognition (clustering) techniques to the workload characterization problem. They allowed problems of higher dimensions to be solved and eliminated the dependence on the analyst's interpretive skill and judgment. Agrawala and Mohr reported on their initial studies of Univac workloads at the University of Maryland (AGR76) and Artis reported on studies of IBM workloads at Bell Laboratories (ART76). Since that time, these and a number of other authors have published numerous papers on the subject in various publications, such as the Computer Measurement Group (CMG) proceedings.

Data Clustering analysis tool allows you to apply cluster-based techniques to the CA MICS database. The software can be applied to any of the unsummarized files that are supported in the DETAIL timespan. The program identifies clusters (that is, similar patterns of activity) in the workload using the SAS FASTCLUS procedure.

The use of clustering techniques introduces other statistical

issues. Among these are scaling and the treatment of outliers. Improper scaling and failure to account for the influence of abnormally large observations are the most common reasons for invalid workload characterization studies.

3.3.2 Usage Guidelines

Before the introduction of cluster-based workload characterization, a number of other workload characterization techniques were utilized. These approaches provide valuable insight into potential errors in the workload characterization process.

The equivalent job approach fails to provide an adequate characterization, since it lacks the important property of robustness. Relatively small changes in the workloads can cause large changes in the definition of an equivalent job. This means that the results of the equivalent job approach are too closely associated with the data used in the analysis to provide general insight into the workloads on the system. Cluster-based workload characterization can replace commonly practiced ad hoc techniques.

The most important issues regarding implementation concern the type of data elements that are used to characterize the clusters and the treatment of unusual points that cannot be represented by clusters. These points, known as outliers, can have a profound effect on the clustering study.

These topics and the use of certain CA MICS database files for clustering analysis are discussed in the following sections:

- 1 - Historical Approaches
- 2 - Basic Hypotheses
- 3 - Application of Workload Characterization
- 4 - Data Selection and Outliers
- 5 - Suggested CA MICS Input Files

3.3.2.1 Historical Approaches

One early approach presented to performance analysis was the equivalent job methodology presented by Jenkins in 1977 (JEN77). Most simply described, the resources consumed by all of the workload elements (for example, job, steps, transactions) over some number of weeks or months are divided by the number of workload elements that were executed during the period. The result of this division is called an equivalent workload element.

An equivalent workload concept initially appears to be very attractive since you can estimate future resource requirements by simply estimating the number of equivalent workload elements. This value is multiplied by the equivalent workload description to calculate the future resource requirements.

3.3.2.2 Basic Hypotheses

The problems associated with the equivalent workload element approach can best be introduced by stating the two hypotheses that were introduced by Artis for testing proposed workload characterizations (ART76). These hypotheses are the "representative hypothesis" and the "stationary hypothesis."

The representative hypothesis tests how well the proposed characterization represents each element of the workload. Rather than testing the proximity of the mean value of the workload to the proposed characterization, you must evaluate the differences between the proposed characterization and each element of the workload.

For example, consider the pathological example of a system that runs 100 testing jobs using less than 5 CPU seconds and 100 production jobs using about 10 CPU minutes each day. (Of course, real workloads are much more complex.) This would result in an equivalent job that required about 5 CPU minutes. Unfortunately, no job in the workload has resource requirements similar to this equivalent job. Therefore, rather than simplifying the problem by reducing the number of workload elements to be considered, we have changed the nature of the problem.

The characteristics of a system that runs 200 equivalent (that is, average) jobs per day would be much different from those of the hypothetical system discussed above. It is interesting to note that the equivalent workload element can only hope to pass the representative hypothesis for systems where all of the workload elements are approximately equal. To be valid, any proposed workload characterization must pass the representative hypothesis test.

The stationary hypothesis evaluates the likelihood that the proposed workload characterization is valid for future workloads. To test this hypothesis, you must wait and test the characterization's applicability over a number of weeks or months. (Alternatively, you may produce the characterization on older historical data and compare it to more recent history that is not being used to do the characterization.)

You should never consider a workload characterization as a tool for forecasting unless this hypothesis has been tested. Using the hypothetical system discussed above, the smallest change in the ratio of test jobs to production jobs would significantly change the workload characterization. Moreover,

even small changes in the average resource consumptions for testing or production jobs would have the same effect. Therefore, it is unlikely that any equivalent workload element approach would ever pass the stationary hypothesis test.

In general, very few proposed workload characterizations will ever pass the stationary hypothesis test. This does not mean that most workload characterizations are invalid; it does mean that most are unsuitable for use as forecasting units.

3.3.2.3 Application of Workload Characterization

Perhaps one of the most common workload characterization studies is the determination of job classes. Typically, the analyst guesses at a solution and then evaluates how well it fits the system's workload. This is known as the ad hoc approach. For example, consider a hypothetical system where the following job class structure has been proposed.

JOB CLASS	CPU MINS	PRINT LINES
=====	=====	=====
A	2	5,000
B	5	20,000
C	unlim	unlim

If 60 percent of the jobs are assigned to class A, 25 percent to class B, and 15 percent to class C, the analyst might assume that the classes provide a good representation of the workload.

Applying the representative hypothesis test could reveal problems. For example, if 30 percent of the jobs use less than 5 CPU seconds and print less than 2,000 lines, then they would be poorly represented by the class A limits. Moreover, their service characteristics will be much poorer when they are randomly mixed with the larger jobs in class A than they will be if they are assigned to a class by themselves.

Another problem resulting from arbitrarily established job classes is the specification of a job class limit that bisects a natural structure in the data. For example, if 15 percent of the jobs normally require between 110 and 130 CPU seconds, imposing a two-minute limit could force the users of jobs into class B to ensure that they are not canceled at the two-minute limit. This problem is also identified by the representative hypothesis test since the approximate two CPU minute jobs are poorly represented by the class B limits.

The characteristics of an installation's job mix are continually evolving. As a result, it is probably unreasonable to apply the stationary hypothesis test to a job class structure. However, you should examine the job class structure on a quarterly or semiannual basis.

Both the equivalent job and ad hoc batch approaches to

workload characterization discussed here show the problems that result from using mean values of arbitrarily selected limits. To meet the requirements of the first hypothesis, we must exploit the natural structure of the workload. Statistical pattern recognition techniques (clustering) offer a powerful tool for determining whether a natural structure exists in an installation's workload data.

3.3.2.4 Data Selection and Outliers

Selecting data elements for the clustering procedure is one of the most important steps in the workload characterization process. In clustering terminology, these data elements are known as features. The result of the clustering process is a ranking of similar types of workloads, based on the mean values of these features.

The objective is to select the minimum number of features that can be used to accurately describe the characteristics of the workload to be analyzed. A small number of features is desirable, since it greatly reduces the number of sources of variability in the characterization, making the results more robust and easier to understand. Note that each time an additional feature is added, splits may occur in as many as half of the existing clusters.

Another consideration in the feature selection process is the representation of outliers. Outliers are data elements that cannot be generalized as similar to other data elements in any of the chosen features. One example of an outlier is a batch job which falls into a CPU-bound loop due to a programming error. Another example is a periodic process, such as a monthly close of a financial application, which is run rarely but requires a large number of resources. Outliers have the highest probability of being poorly represented by the developed workload characterization.

There are two classes of features that are normally used in workload characterization studies:

- o Rate-based features, such as EXCPs per CPU second
- o Total features, such as CPU minutes

Rate-based features are desirable because they allow you to control the content of the outliers. For example, consider the use of CPU minutes and EXCPs as total features for characterizing the resource consumption of a workload. Because the largest jobs would be outliers, the resulting characterization would poorly represent them. Unfortunately, these jobs might represent a significant percentage of the resources that are consumed.

To avoid this problem, you could use rate-based features. When you divide EXCPs by CPU seconds, you change the content of the outliers. Rather than representing jobs that consume significant resources, they represent jobs with exceptional EXCP/CPU second rates. The probability that these jobs would

consume significant resources is very small, since the ratio of elapsed time to CPU resource consumption for I/O-bound jobs is very high.

Such rate-based features have been called instantaneous workload descriptions, since they characterize the activities of a workload element for an instant of CPU time (ART76, ART78).

Although using total features introduces concerns about representing outliers, they have many applications to capacity planning and performance problems. For example, using total features for job and transaction class structures treats the occasional massive workload element as an exception. By handling the occasional element as an exception, the system can concentrate on exploiting the characteristics of the majority of the workload elements to simplify scheduling. You must, however, always be aware of the content of the outliers and pay particular attention to the marginally assigned resource vectors to verify that the objectives of the study are not compromised.

3.3.2.5 Suggested CA MICS Input Files

Although you may analyze any of the unsummarized CA MICS files in the DETAIL timespan, a number of files may particularly interest you. In the Batch Information Area, they are as follows:

BATJOB	Contains one record for each job
BATPGM	Contains one record for each job step, that is, program
BAT_TS	Contains one record for each TSO session
BAT_ST	Contains one record for each started task

In the TSO Information Area, the following files may interest you:

TS0TSU	Contains one record for each TSO/MON interval in which a user was active
TS0TSI	Contains one record for each command or program executed interactively

In the CICS Information Area, you can create files in the DETAIL timespan that contain one record for each transaction using the _USRXfff macros. These macros are discussed in Chapter 10 of the CA MICS Analyzer Option for CICS Guide. In the event that you use this option, the following files may interest you:

CICCSU	Contains one record for each transaction sorted by user
CICCAU	Contains one record for each transaction sorted by application

You can also create IMS transaction files in the DETAIL timespan that contain one record for each transaction using the _USRXfff macros. These macros are discussed in Chapter 10 of the CA MICS Analyzer Option for IMS Guide. In the event that you choose this option, the following files may interest you:

IMISISU	Contains one record for each transaction sorted by user
IMSIAU	Contains one record for each transaction sorted

by application

Although you can analyze other CA MICS information areas and files using Workload Characterization, no results have been published that detail the application of clustering algorithms to other than workload element data.

3.3.3 Standard Output

Data Clustering using Neugents technology produces reports and simple graphs that describe the original data, show the clusters that are identified from this data, and analyze the degree of fit with which these clusters characterize the data.

The 5 reports and 2 graphs are:

- o Clustering Input Descriptive Statistics. This report details the characteristics of the variables that are specified as criteria for the clustering analysis.
- o Cluster Performance Summary. This report consists of three sections: Clustering Execution Options, a summary of the options selected by the analyst for the study; Cluster Feature Contents, a list of each of elements or features used for clustering, and the boundaries selected for determining normalcy of the data values; and finally, the Cluster Population Summary, which lists each cluster, its size, its percent of population, and an index used to indicate relative performance of the cluster.
- o Cluster Population Summary. This report provides the resource consumption characteristics of each of the clusters. For each cluster, its size and population are summarized and all clustering and reporting elements are listed, along with their respective usage values. This report permits you identify resource consumption patterns that may represent performance or capacity problems.
- o Data Value Exception Detail. This report provides a detail listing of any observation that has one or more clustering elements that are identified as outliers. The report is sequenced in the same order as is the original CA MICS files used as input and can be used to further investigate why certain cases have been identified as outliers and whether or not they actually represent abnormal situations.
- o Sparse Cluster Population. This report provides a listing of all clusters that have been identified as 'sparse' and their population contents. A sparse cluster is defined within this application as one having less than 0.5% of the overall population of the study. For example, if a sample of 2,000 cases were selected for analysis, then a sparse cluster would contain 10 or fewer observations.

- o Cluster Index Analysis. This graph depicts a measurement of relative performance for each cluster, by presenting an index calculated by Neugents technology in terms of the root mean square of the distances of the resource vectors or feature values from the cluster centroid. The cluster index of each cluster is displayed in either a Block or a horizontal bar chart, depending on the number of clusters being presented. This chart permits the analyst to quickly access the overall performance of the workload characterization effort.
- o Cluster Population Analysis. This graph depicts the population of the top "n" clusters in a PIE chart, which can be used to understand pictorially how the study data actually populated the clusters.

The following sections describe the fields, data elements, and format of each report.

- 1 - Clustering Input Descriptive Statistics Report
- 2 - Cluster Performance Summary Report
- 3 - Cluster Population Summary Report
- 4 - Data Value Exception Detail Report
- 5 - Sparse Cluster Population Report
- 6 - Cluster Index Analysis Graph
- 7 - Cluster Population Analysis Graph

3.3.3.1 Clustering Input Descriptive Statistics Report

The Descriptive Statistics Report provides a summary of the statistics that are calculated for the cluster features. Figure 3-33 illustrates a sample Descriptive Statistics Report.

Data Clustering Analysis					
Clustering Input Descriptive Statistics					
For: Thursday, June 19, 2003					
Summary for Total Sample:					
Number of Sample Observations:		2,000			
Feature	Minimum	Maximum	Average	Std. Dev.	CV
JOBTCBTM	0:00:00.01	0:40:00.28	0:00:13.14	0:01:42.89	7.83
JOBEDASD	1.00	1,154,523.00	5,992.30	34,582.08	5.77
Summary for Trimmed Sample:					
Number of Sample Observations:		1,919			
Feature	Minimum	Maximum	Average	Std. Dev.	CV
JOBTCBTM	0:00:00.01	0:01:00.81	0:00:03.74	0:00:07.71	2.06
JOBEDASD	1.00	38,096.00	2,505.00	4,508.48	1.80

Figure 3-33. Descriptive Statistics Report

Note that the report has two sections: one for the original sample randomly taken from the input data; and the other, the sample data after trimming, using the Sample Trim Limit value that you specify on the Clustering Execution Parameters Option screen. See Section 3.3.5.9.

The Descriptive Statistics Report contains the following fields:

OBSERVATIONS: The number of observations selected for processing.

FEATURE: The feature name. The feature names that are listed in the report correspond to the features that are specified on the Data Clustering screen. In Neugents technology, this is also called a "pattern".

MINIMUM: The minimum value observed for the feature in the sample.

MAXIMUM:	The maximum value observed for the feature in the sample.
AVERAGE:	The average calculated for the feature observations in the sample.
STD DEV:	The standard deviation calculated for the feature observations in the sample.
CV:	The coefficient of variation. The CV is calculated by dividing the standard deviation by the average.

3.3.3.2 Cluster Performance Summary Report

A sample Cluster Performance Summary report is shown below in Figure 3-34.

```

Data Clustering Analysis                                1
Cluster Performance Summary
For: Monday, June 23, yyyy

Clustering Execution Options

Cluster Input Data: SAMPLE
Clustering Method: RADIUS
Maximum Cluster Radius: 3.0

Training Sample Size: 2000
Sample Trim Limit: 97.5 Pct.
Sampling Percentage : 2.5 Pct.

Include Outliers in Clusters: YES
Include Sparse Clusters: YES
Sparse Cluster Limit: 0.5 Pct.

Report Cluster Population: YES
Report Using Account Codes: YES
Report Sparse Clusters : YES
Report Outliers Separately: YES
Create Cluster Index Graph: YES
Create Population Graph: YES

Input CA MICS Files: BATJOB
Input Dataset Name: 'COGDA01.AUDIT.BATJOB'
-----

Cluster Feature Contents

Feature Name ----- Description ----- Lower Upper Cases
-----
JOBTCBTM Job TCB CPU Time 0:00:00.00 0:01:05.70 50
JOBEDASD DASD EXCPS 0.00 38,933.00 50
=====
100

Note: A given case may be classified as an outlier more than once if multiple
analysis elements contain abnormal values. Therefore, the totals in this report
section will not necessarily agree with those of the Population section
-----

```

Data Clustering Analysis
Cluster Performance Summary
For: Monday, June 23, yyyy

Cluster Population Summary

Cluster	Radius	Normal Obs.	Outlying Obs.	Total Obs.	% of Population	Clustering Index
-----	-----	-----	-----	-----	-----	-----
1	0.00	0	14	14	0.70	0.00
2*	0.82	0	2	2	0.10	0.82
3*	0.82	0	6	6	0.30	0.51
4*	1.17	0	2	2	0.10	1.17
5*	1.20	0	2	2	0.10	1.20
6*	1.39	4	0	4	0.20	1.01
7*	1.43	0	3	3	0.15	1.02
8*	1.50	0	2	2	0.10	1.50
9*	1.55	0	5	5	0.25	1.22
10	1.61	1,541	0	1,541	77.05	0.34
11	1.71	15	0	15	0.75	0.77
12*	1.86	0	9	9	0.45	1.26
13*	2.01	0	2	2	0.10	2.01
14	2.04	4	8	12	0.60	1.19
15	2.07	129	0	129	6.45	0.95
16	2.22	12	0	12	0.60	0.80
17	2.25	0	11	11	0.55	1.69
18*	2.29	0	2	2	0.10	2.29
19*	2.45	0	2	2	0.10	2.45
20*	2.54	0	3	3	0.15	1.99
21*	2.70	0	2	2	0.10	2.70
22*	2.71	2	6	8	0.40	1.68
23	2.72	39	0	39	1.95	1.70
24	2.83	173	0	173	8.65	0.98
		=====	=====	=====	=====	
		1,919	81	2,000	100.00	

'*' Indicates that the cluster is sparsely populated.

Sparse cluster are defined as having a population that is less than 0.5 % of the total population. In this study, the sparse cluster population limit is 10 cases.

Figure 3-34. Cluster Performance Summary report

The Cluster Performance Summary report is presented in three sections as described next:

The first section explains the execution options chosen by the analyst for the study being performed and serves to document which reports, graphs will be produced.

For an explanation of the executions options and their respective usage, see section 15.x of the product guide.

The second section of the report presents the feature contents of the clusters, including the limits used to determine outliers and the number of times (cases) observations were declared as outliers because of a given feature value.

FEATURE NAME:	The name of the data element chosen for clustering. This is normally a CA MICS data element but can be a computed (user defined) element if required.
DESCRIPTION:	The SAS label of the selected data element, from either the CA MICS GENLIB definition of supplied by the user.
OUTLIER LIMITS:	Statistical bounds to determine if a given data value is considered "normal" or represents an abnormal condition. These values are computed based on the Sample Trim Limit value specified on the Clustering Execution Options screen.
LOWER:	Cases (observations) whose value for this feature are less than this value are defined as outliers. Given that most performance measurement data is positive in scope, the current implementation uses a value of zero for this boundary.
HIGHER:	Cases (observations) whose value for this feature are greater than this value are defined as outliers.
CASES:	The number of cases (observations) that were declared as outliers because of this feature value. Note that a given case may be flagged for multiple times and therefore the number of cases presented in this section may exceed the total number of Outlying observations in the

next report section.

A footnote to this affect is printed at the bottom of this report section.

The third report section presents a summary of the size, population and general performance of each cluster.

CLUSTER NUMBER:	The cluster number. Cluster numbers are assigned sequentially to the patterns that are identified by the algorithm. Note that the order in which the clusters are identified is not an indicator of merit.
RADIUS:	The geometric distance from the cluster center centroid to the outer boundary of the cluster, expressed in terms of Standard Deviations. You define the outer limit of this value by supplying the Maximum Cluster Radius value on the Clustering Execution Options screen.
NORMAL OBS:	The count of observations within this cluster where all feature values were found to be "normal" in a statistical sense. In this implementation, normal feature values are those that are less than the value determined by the Sample Trim Limit. For example, if the Sample Trim Limit of 97.5%, then all feature values of "normal" clusters would reside below the 97.5 percentile of the sample.
OUTLYING OBS:	The count of observations within this cluster where one or more feature values were found to be "outliers" in a statistical sense. In this implementation, outlying feature values are those that are greater than the value determined by the Sample Trim Limit. For example, if the Sample Trim Limit of 97.5%, then at least one feature value of "outlying" clusters would reside above the 97.5 percentile of the sample.
TOTAL OBS:	The sum of NORMAL and OUTLYING observations.
% OF POPULATION:	The percent of the total sample populations represented by this cluster.
CLUSTERING	Formally called the Performance Index, this

INDEX: metric was renamed in this implementation to avoid confusion with similar terms in the z/OS Workload Manager. It is the Root Mean Square of the distances of all cases (observations) within the cluster and serves as a simple measure of clustering effectiveness. The lower this value, the tighter the fit of the cluster data. This is not to say that outliers are not present; only that they are not distorting the cluster shape by their presence.

Note the footnote at the bottom of this report section. It refers to "sparse" clusters and presents the definition and limits used for determining which clusters are considered "sparse". Sparse clusters are generally populated by outliers and are often dropped from further analysis after they have been reviewed.

3.3.3.3 Cluster Population Summary Report

A sample Cluster Population Summary report is shown below in Figure 3-35.

Data Clustering Analysis							1
Cluster Population Summary							
For: Tuesday, June 24, yyyy							
----- Cluster: 1 -----							
Radius: 0.00 Clustering Index: 0.00 Maximum Cluster Radius: 3.0 STDs							
Normal Obs.:	0	0.00%	Observations << 97.5	Pct.			
Outlying Obs.:	14	0.70%	Observations >> 97.5	Pct.			
Total Obs.:	14	0.70%					
Cluster Feature Resources							
Feature	Description	Minimum Value	Average Value	Standard Deviation	Maximum Value	Total Value % of Total	
JOBTCBTM	Job TCB CPU Time	25	656	687	2,400	9,186 34.96	
JOBEDASD	DASD EXCPS	5,103	236,581	298,042	1,154,523	3,312,131 27.64	
Cluster Report Resources							
Feature	Description					Total Value % of Total	
JOBNLR	Total Logical Writer Records					0 0.00	
----- Cluster: 10 -----							
Radius: 1.61 Clustering Index: 0.34 Maximum Cluster Radius: 3.0 STDs							
Normal Obs.:	1,541	77.05%	Observations << 97.5	Pct.			
Outlying Obs.:	0	0.00%	Observations >> 97.5	Pct.			
Total Obs.:	1,541	77.05%					
Cluster Feature Resources							
Feature	Description	Minimum Value	Average Value	Standard Deviation	Maximum Value	Total Value % of Total	
JOBTCBTM	Job TCB CPU Time	0	1	2	14	1,873 7.13	
JOBEDASD	DASD EXCPS	1	914	1,155	6,523	1,408,821 11.76	
Cluster Report Resources							
Feature	Description					Total Value % of Total	
JOBNLR	Total Logical Writer Records					19,968 99.56	
Sparse Clusters Have Been Excluded							

Figure 3-35. Cluster Population Summary Report

Data Clustering Analysis
Cluster Population Summary

For: Monday, June 23, yyyy

The following clusters were determined to be 'sparse' in their population and have been excluded from processing.

Cluster:	Population
2	2
3	6
4	2
5	2
6	4
7	3
8	2
9	5
12	9
13	2
18	2
19	2
20	3
21	2
22	8

CLUSTER NUMBER:	Cluster numbers are assigned sequentially to the patterns that are identified by the algorithm. Note that the order in which the clusters are identified is not an indicator of merit.
RADIUS:	The geometric distance from the cluster center centroid to the outer boundary of the cluster, expressed in terms of Standard Deviations. The outer limit of this value is defined by the user as the Maximum Cluster Radius value on the Clustering Execution Parameters screen.
CLUSTERING INDEX:	Formally called the Performance Index, this metric was renamed in this implementation to avoid confusion with similar terms in the z/OS Workload Manager. It is the Root Mean Square of the distances of all cases (observations) within the cluster and serves as a simple measure of clustering effectiveness. The lower this value, the tighter the fit of the cluster data. This is not to say that outliers are not

present; only that they are not distorting the cluster shape by their presence.

MAXIMUM CLUSTER RADIUS:	The maximum size of any cluster in the study under consideration, expressed in terms of standard deviations. This value is specified by the analyst on the Execution Options for Neugents Technology panel (PERG900U) and is presented here for documenting the cluster definitions.
NORMAL OBS:	The count of observations within this cluster where all feature values were found to be "normal" in a statistical sense. In this implementation, normal feature values are those that are less than the value determined by the Sample Trim Limit. For example, if the Sample Trim Limit of 97.5%, then all feature values of "normal" clusters would reside below the 97.5 percentile of the sample.
OUTLYING OBS:	The count of observations within this cluster where one or more feature values were found to be "outliers" in a statistical sense. In this implementation, outlying feature values are those that are greater than the value determined by the Sample Trim Limit. For example, if the Sample Trim Limit of 97.5%, then at least one feature value of "outlying" clusters would reside above the 97.5 percentile of the sample.
TOTAL OBS:	The sum of NORMAL and OUTLYING observations.

CLUSTER FEATURE RESOURCES:

For each feature (clustering element) that you define on the Clustering Executions Options screen, a separate line will be generated under the Cluster Feature Resources entry on the report, and will contain the following elements:

FEATURE:	The name of the data element chosen for clustering. This is normally a CA MICS data element but can be a computed (user defined) element if required.
DESCRIPTION:	The SAS label of the selected data element, from either the CA MICS GENLIB definition of

supplied by the user.

MINIMUM VALUE:	The minimum of all values for this feature within this cluster.
AVERAGE VALUE:	The average of all values for this feature within this cluster. This value approximates the centroid value for this feature and is also referred to as the feature mean.
STANDARD DEVIATION:	The standard deviation of all values for this feature within this cluster.
MAXIMUM VALUE:	The maximum of all values for this feature within this cluster. This value approximates the outer boundary value for this feature.
TOTAL VALUE:	The sum of all values for this feature within this cluster.
% OF TOTAL:	The percentage of the sum of this feature's values for this cluster compared to the sum for the entire sample population. For example, in the report above, the JOBTCBTM represented by cluster 1 is nearly 35% of the JOBTCBTM for the entire sample.

\$SCOL

CLUSTER REPORT RESOURCES:

For each feature (reporting element) that you define on the Clustering Executions Options screen, a separate line will be generated under the Cluster Report Resources entry on the report, and will contain the following elements:

FEATURE:	The name of the data element chosen for clustering. This is normally a CA MICS data element but can be a computed (user defined) element if required.
DESCRIPTION:	The SAS label of the selected data element, from either the CA MICS GENLIB definition of supplied by the user.
TOTAL VALUE:	The sum of all values for this feature within this cluster. This value represents the total of this reporting element for this cluster.
% OF	The percentage of the sum of this feature's

TOTAL: values for this cluster compared to the sum for the entire sample population. For example, in the report above, the JOBNLR represented by cluster 1 is 0% of the JOBNLR for the entire sample.

If the INCLUDE SPARSE CLUSTERS option on the clustering Execution Options screen has been set to "Yes", a page of the report will contain a listing of all excluded clusters and their populations.

The following fields are presented in this section:

**CLUSTER
NUMBER:** Cluster numbers are assigned sequentially to the patterns that are identified by the algorithm. Note that the order in which the clusters are identified is not an indicator of merit.

POPULATION: The population of each excluded cluster.

3.3.3.4 Data Value Exception Detail Report

This is a simple listing of any observation that contains one or more outlying data values.

The application will determine the key sequence structure of the input file, modify the key structure to remove date and time key values except for the End Timestamp, and present the listing using the modified key sequence. You can optionally exclude the CA MICS account codes from the reports by specifying a value of "NO" to the REPORT USING ACCOUNT FIELDS option on the Clustering Execution Options screen. This report can be used to verify that the outliers are truly exceptional, and not merely unusual workloads.

Note that each outlying clustering data value is flagged with an "*". This allows you to identify outliers in the data and understand which elements(s) contained abnormal data values.

A sample Data Value Exception Detail report is shown below in Figure 3-36.

3.3 Data Clustering Analysis

Data Clustering Analysis
Data Value Exception Detail
For: Monday, June 23, yyyy

Listing of Observations With Outlier Values

System Identifier	DIVISION	DEPARTMENT	RESERVED	Job Group	Job Identification	Job TCB CPU End Time Stamp	Time	DASD EXCPS
XAL1	BNK DEV	01622 01056 01060 01062	*****	7	FCSBK150	07NOVyy:10:08:53.48	0:01:49.84*	65,053*
					LTSIPCSX	14NOVyy:17:54:24.33	0:40:00.28*	354,840*
					CYGIPCS	12NOVyy:10:46:57.98	0:01:41.88*	178,301*
					BEAMA05M	07NOVyy:13:30:34.99	0:00:06.64	90,943*
					IQBAM02D	06NOVyy:13:50:11.44	0:02:10.87*	34,038
					JONCA04B	12NOVyy:15:43:33.43	0:03:32.57*	2,879
						12NOVyy:15:55:19.91	0:01:10.59*	2,702
						12NOVyy:16:10:50.76	0:01:14.38*	2,661
					JONCA043	14NOVyy:12:03:07.41	0:00:24.55	126,664*
					DIXPA02D	15NOVyy:10:46:48.50	0:01:17.18*	15,968
						15NOVyy:13:30:36.48	0:01:14.50*	15,781
					JSIEXTR	11NOVyy:14:33:57.06	0:00:03.56	48,436*
					WYKCH01D	11NOVyy:15:27:19.09	0:02:15.77*	25,010
						12NOVyy:10:24:31.11	0:01:23.27*	26,935
					MANDA06L	04NOVyy:14:41:52.90	0:00:16.10	46,024*
					MANDA06M	04NOVyy:14:42:19.03	0:00:16.12	46,024*
					JERMI01L	11NOVyy:15:31:18.24	0:00:02.36	42,283*
						11NOVyy:15:34:53.13	0:00:02.16	39,770*
					JERMI01M	11NOVyy:17:08:46.41	0:00:02.49	44,844*
						11NOVyy:17:16:02.77	0:00:02.85	51,817*
						11NOVyy:17:25:56.11	0:00:02.38	44,865*
						11NOVyy:17:46:16.77	0:00:02.12	42,335*
						11NOVyy:17:49:34.96	0:00:04.81	101,604*
						11NOVyy:17:51:30.57	0:00:02.46	48,545*
						11NOVyy:17:53:40.22	0:00:04.31	89,564*
						11NOVyy:17:55:53.45	0:00:02.67	53,444*
						12NOVyy:12:58:13.85	0:00:02.91	50,464*
					FEESE01Y	07NOVyy:15:09:18.12	0:00:31.65	43,576*
						07NOVyy:15:34:08.89	0:00:59.88	79,414*
						15NOVyy:11:34:19.60	0:00:33.91	45,279*
						15NOVyy:11:00:52.65	0:00:05.13	95,660*
GIS	01266 00045	*****	7	0	HENR004Q	12NOVyy:11:11:59.69	0:01:31.50*	9,099
					A02ICST3	08NOVyy:17:22:57.44	0:01:37.44*	5,808
					MPIRSADT	11NOVyy:11:23:34.38	0:00:17.46	99,448*
					TCWFAST\$	11NOVyy:11:50:11.47	0:00:19.64	156,990*
						14NOVyy:11:47:28.60	0:03:38.29*	10,763
					TESTXCBK	06NOVyy:17:18:56.79	0:07:28.57*	47,494*
					TOPCDEV1	08NOVyy:09:51:53.06	0:05:40.66*	148,789*
						08NOVyy:11:42:22.55	0:01:40.63*	7,267
						08NOVyy:13:41:37.00	0:01:38.81*	7,179
						08NOVyy:13:49:04.83	0:01:33.66*	7,094
					TOPCQAQA	12NOVyy:10:05:58.65	0:03:00.82*	14,794
						14NOVyy:18:24:00.78	0:01:44.87*	25,197
					TOPDDEV2	13NOVyy:10:25:47.61	0:20:52.34*	538,224*
					TOPDQAQA	06NOVyy:11:35:54.43	0:17:50.60*	1,154,523*
					WKBKVM82	04NOVyy:12:20:36.27	0:21:08.57*	225,024*
					WKBKVM83	04NOVyy:18:11:15.80	0:23:50.50*	230,790*
					WKBKVM84	11NOVyy:10:52:19.66	0:20:15.46*	220,572*
					WKBKVM85	11NOVyy:12:50:10.33	0:25:27.51*	220,572*

Note: '*' Indicates that the data value is an outlier

Figure 3-36. Data Value Exception Detail report

The Data Value Exception report is dynamically built and does not have a fixed format. The first 'n' columns of the report contain the sequence fields of the CA MICS file used for analysis. These fields are documented in the CA MICS Data Dictionary. The remaining elements are the ones that you specify on the Data Clustering screen as clustering elements. If these are CA MICS data elements, which is the typical case, then their definitions and usage data can be located in the CA MICS Data Dictionary as well.

3.3.3.5 Sparse Cluster Population Report

This is a listing of all clusters that have been defined as "sparse". Sparse clusters are by definition those clusters having less than 0.5% of the population of the sample. Generally, these clusters contain outliers but may also contain normal cases.

The application will determine the key sequence structure of the input file, modify the key structure to remove date and time key values except for the End Timestamp, and present the listing using the modified key sequence. You can optionally exclude the CA MICS account codes from the reports by specifying a value of "NO" on the REPORT USING ACCOUNT FIELDS option on the Clustering Execution Options screen. This report can be used to verify that the sparse clusters are truly exceptional cases, not merely unusual cases.

Note that each outlying clustering data value is flagged with an "*". This allows you to identify outliers in the data, and understand which elements(s) contain abnormal data values.

A sample Sparse Cluster Population report is shown below in Figure 3-37.

Sparse Cluster Population Report									
For: Monday, June 23, yyyy									
Listing of Observations Within Cluster									
Cluster Number	System Identifier	DIVISION	DEPARTMENT	RESERVED	Job Group	Job Identification	End Time Stamp	Job TCB CPU Time	DASD EXCPS
2	XAL1	GIS	00045	*****	5	MPGLAUDT	07NOV02:11:39:02.05	0:02:06.60*	7,340
	XAL1	GIS	02046	*****	0	T1PR2NEW	13NOV02:17:37:26.57	0:02:09.01*	108
3	XAL1	GIS	00045	*****	0	A02ICST3	12NOV02:11:11:59.69	0:01:31.50*	9,099
	XAL1	GIS	00045	*****	0	MPIRSADT	08NOV02:17:22:57.44	0:01:37.44*	5,808
	XAL1	GIS	00045	*****	0	TOPCDEV2	08NOV02:11:42:22.55	0:01:40.63*	7,267
	XAL1	GIS	00045	*****	0	TOPCDEV2	08NOV02:13:49:04.83	0:01:33.66*	7,094
	XAL1	UNK	UNK	*****	0	CACMPQA	14NOV02:10:09:47.07	0:01:41.14*	8,383
4	XAL1	GIS	00045	*****	5	MPMINAS	14NOV02:16:45:15.65	0:00:35.54	70,439*
	XAL1	GIS	00045	*****	5	MPMIPACE	14NOV02:10:30:56.85	0:00:40.26	80,646*
5	XAL1	DEV	01076	*****	7	FEESE01Y	07NOV02:15:34:08.89	0:00:59.88	79,414*
	XAL1	GIS	02045	*****	0	BEHMAMDB	11NOV02:15:08:03.50	0:01:18.45*	79,294*
6	XAL1	DEV	01062	*****	5	MCCTI03H	07NOV02:14:32:14.63	0:00:34.83	23,429
	XAL1	DEV	01062	*****	7	DIXPA02M	15NOV02:15:35:10.77	0:00:42.85	26,409
	XAL1	GIS	02045	*****	0	BEHMAMDB	11NOV02:15:31:53.21	0:00:27.57	27,384
	XAL1	GIS	02045	*****	0	BEHMAMEN	11NOV02:14:52:19.61	0:00:24.20	23,585
7	XAL1	GIS	00045	*****	0	TCWFAST\$	11NOV02:11:50:11.47	0:00:19.64	156,990*
	XAL1	GIS	00045	*****	5	STARLOAD	13NOV02:15:58:32.57	0:00:28.49	148,507*
	XAL1	GIS	00045	*****	5	STARLOAD	13NOV02:16:26:38.42	0:00:26.74	148,503*
8	XAL1	GIS	00045	*****	0	WKBKVM85	11NOV02:12:50:10.33	0:25:27.51*	220,572*
	XAL1	GIS	02046	*****	7	MELR101B	07NOV02:18:12:06.64	0:25:34.91*	233,415*
9	XAL1	DEV	01062	*****	0	BEAMA05M	07NOV02:13:30:34.99	0:00:06.64	90,943*
	XAL1	DEV	01076	*****	0	JERMI01M	11NOV02:17:49:34.96	0:00:04.81	101,604*
	XAL1	DEV	01076	*****	0	JERMI01M	11NOV02:17:53:40.22	0:00:04.31	89,564*
	XAL1	DEV	01266	*****	7	HENR004Q	15NOV02:11:00:52.65	0:00:05.13	95,660*
	XAL1	GIS	00045	*****	0	TCWFAST\$	11NOV02:11:23:34.38	0:00:17.46	99,448*
12	XAL1	DEV	01062	*****	5	JONCA04B	12NOV02:15:55:19.91	0:01:10.59*	2,702
	XAL1	DEV	01062	*****	5	JONCA04B	12NOV02:16:10:50.76	0:01:14.38*	2,661
	XAL1	DEV	01062	*****	7	DIXPA02D	15NOV02:10:46:48.50	0:01:17.18*	15,968
	XAL1	GIS	00045	*****	5	VRPDIV	08NOV02:17:27:14.42	0:01:23.85*	9,966
	XAL1	GIS	02046	*****	0	FAIRSCAN	07NOV02:11:39:42.42	0:01:14.27*	13,231
	XAL1	GIS	02046	*****	0	FAIRSCAN	08NOV02:14:40:26.10	0:01:13.71*	13,234
	XAL1	UNK	UNK	*****	0	DBTNGFR2	14NOV02:09:26:28.85	0:01:12.29*	5,774
	XAL1	UNK	UNK	*****	0	T1CAPS2	11NOV02:09:07:14.75	0:01:21.77*	15,217
13	XAL1	DEV	01060	*****	5	CYGIPCS	12NOV02:10:46:57.98	0:01:41.88*	178,301*
	XAL1	SPT	04733	*****	5	BEEST02U	08NOV02:11:58:26.55	0:02:11.93*	173,919*

Note: '*' Indicates that the data value is an outlier

Figure 3-37. Sparse Cluster Population report

The Sparse Cluster Population report is dynamically built, and does not have a fixed format. The first 'n' columns of the report contain the sequence fields of the CA MICS file used for analysis. These fields are documented in the CA MICS Data Dictionary. You specify the remaining elements on the Data Clustering screen as clustering elements. If these are CA MICS data elements, which is the typical case, then their definitions and usage data can be found in the CA MICS Data Dictionary as well.

Figure 3-38. Cluster Index Analysis graph

The clustering index value is the Root Mean Square (RMS) of the distances of all feature values within the cluster. As stated previously, this is a relative measurement of cluster performance and "goodness of fit" of the data, with respect to the cluster centroid values. A lower value represents a better performing cluster, indicating more tightly formed clusters and fewer outlying values.


This graph is a simple, yet effective to quickly evaluate the overall performance of the clustering operation and can be used to communicate this performance to non-technical personnel.

3.3.3.7 Cluster Population Analysis graph

Cluster Population Analysis is a PIE chart depicting the populations of the top 'n' clusters, where 'n' depends on the printer linesize and pagesize specifications, and the distribution of data values within the sample. For example, the number of 'slices' of the PIE chart may be reduced if one slice is larger than the remaining slices.

Data Clustering Analysis
Cluster Population Analysis
For: Monday, June 23, 2003

Analysis Performed Using Variables: JOBTCBTM JOBEDASD



Sparse Clusters Have Been Excluded

Figure 3-39. Cluster Population Analysis graph

This chart provides a simple, yet effective way of presenting the population distribution of the study data in a high-level summary.

3.3.4 Analytic Technique Tutorial

This section discusses general Data Clustering methodology. While this is not a complete discussion of the topic, it does provide an introduction to the statistical techniques employed by the Neugents technology in this process. The tutorials presented in this section concentrate on the following topics:

- 1 - Clustering Algorithms
- 2 - Scaling
- 3 - Algorithm Implementation

3.3.4.1 Clustering Algorithms

Although there are a wide variety of statistical algorithms that you can use to develop clusters, one of the simplest and most popular algorithms is the K-means algorithm (MAC67). A hybrid of both the hierarchical and the K-means algorithm is employed within Neugents technology for clustering operations.

To understand how clustering algorithms work, you must understand the following definitions.

Resource Vector:	An n-tuple of measurements selected to describe a workload element. For example, you might choose to describe jobs in terms of the three distinct measurements: CPU time, print lines, and tape mounts. Each job would then be represented by a 3-tuple (CPU time, print lines, and tape mounts). This is also sometimes referred to as a usage vector, or simply a vector. Neugents technology uses the term "pattern" when describing such a set of related measurements.
Feature:	A component of a resource vector. In the above example, CPU time, print lines, and tape mounts are the three features of resource vectors describing the workload element, job.
Space:	The n-dimensional coordinate system in which resource vectors are defined.
Feature Mean:	The mean of all values of a given feature taken over a set of resources.
Feature Variance:	The variance (that is, the standard deviations squared) associated with the distribution of values of a given feature, taken over a set of resource vectors. Note that the feature with the greatest variance defines the longest coordinate axis of the space.
Centroid:	The geometric center of the set of all resource vectors contained in a space. The location of the centroid is defined by an n-feature vector. Its component values are the feature means of the resource vectors in the space.
Cluster:	A set of resource vectors in a given space plus

their centroid. Note that the centroid of a set of vectors need not be one of the original vectors in the space.


Geometric Distance:	The distance between any two points (that is, resource vectors) in a given space. It is calculated by squaring the difference between the points (that is, the feature values) for each feature, summing the squares over all features, and taking the square root of the resulting sum. The formula for a two-dimension geometric distance is better known as the Pythagorean theorem. A formula for calculating geometric distance is presented in the following first equation.
Radius:	The greatest geometric distance between a centroid and any of its associated resource vectors.
Relative Position:	The ratio of the geometric distance between a given resource vector or pattern and the radius of the cluster, expressed as a value from 0.00 to 1.00. The lower the value of the relative position, the closer the particular resource vector or pattern is to the cluster center.
Objective Radius:	An externally defined parameter that specifies a maximum desired radius for clusters identified by a clustering algorithm.
Splitting:	The partitioning of a space into two smaller spaces. A space is split in half along the axis of greatest variance. A clustering algorithm will continue to split resource vector spaces as long as the radius of any cluster in the partition is greater than the specified objective radius.

Consider the following two-dimensional clustering example. Fifteen resource vectors were selected for the example. They are shown in the following table:

VECTOR NUM.	X	Y
=====	=====	=====
1	1	1
2	2	1
3	3	1
4	3	8
5	4	7
6	4	8
7	4	9
8	5	8
9	8	4
10	9	3
11	9	4
12	9	5
13	10	4
14	10	8
15	11	8

Note: The range of values selected for the two features in this example are similar. This avoids the problems that are introduced by scaling considerations. Scaling is discussed in Section 15.4.2. The objective radius for this example is 1.

A plot of the resource vectors is shown in Figure 3-40. The resource vectors are represented by asterisks (*) and the centroid of the space is represented by a pound sign (#).






Figure 3-40. Initial Space

The first step in the solution process is to calculate the feature means, variances, and radius for the space. (The centroid is defined by the feature means.) The calculated values are as follows:

SPACE	FEATURE	CENTROID (MEAN)	VARIANCE	RADIUS
all	X	6.13	11.41	6.67
	Y	5.27	8.50	

As mentioned previously, the radius of the space is calculated using the geometric distance equation. This equation states that the distance, D, between the centroid of the k-th cluster, C(k), and the i-th resource vector, R(i), is equal to the square root of the sum of the squares of their feature differences. The general form of the geometric distance equation is shown in Equation 1, where j is the feature number index.

$$D = \sqrt{\sum_{j=1}^n (C(k,j) - R(i,j))^2}$$

(Eqn 1)

For example, the farthest resource vector from the centroid of the space shown in Figure 3-40 is the resource vector at (1,1). The calculation of the distance between the resource vector and the centroid is shown in Equation 2.

$$D = \sqrt{(6.13 - 1.0)^2 + (5.27 - 1.0)^2} \quad (\text{Eqn 2})$$

$$D = \sqrt{26.32 + 18.23} = 6.67$$

As shown in the previous table, the variance in the X direction is 11.4 and the variance in the Y direction is 8.50. Since the radius of the space is greater than the objective radius, the space is split at the centroid along the axis of greatest variance. This results in two spaces, A and B, as is shown in Figure 3-41.




Figure 3-41. First Split

The characteristics of the two spaces resulting from the split are shown in the following table.

SPACE	FEATURE	CENTROID (MEAN)	VARIANCE	RADIUS
A	X	3.25	2.70	4.92
	Y	5.38	13.41	
B	X	9.33	1.07	3.31
	Y	5.14	4.14	

Once again, the radii of the spaces are greater than the objective radius. Hence, each of the spaces is split along the axis of greatest variance. This results in four spaces, A, B, C, and D, as shown in Figure 3-42. Note that in spaces A, C, and D the centroids happen to coincide with a resource vector.




Figure 3-42. Second Split

The characteristics of the four spaces resulting from the split are shown in the following table.

SPACE	FEATURE	CENTROID (MEAN)	VARIANCE	RADIUS
A	X	4.0	0.50	1.0
	Y	8.0	0.50	
B	X	10.5	0.50	0.5
	Y	8.0	0.00	
C	X	2.0	1.00	1.0
	Y	1.0	0.00	
D	X	9.0	0.50	1.0
	Y	4.0	0.50	

Since all of the radii are less than or equal to the objective radii, a solution for the space introduced in Figure 3-42 has been developed. Although this problem is very simplistic, it illustrates the concepts of the K-mean algorithm.

Thus, four clusters have been identified that represent the 15 resource vectors in the space. Although the actual implementation of the clustering algorithm is somewhat more complex, the same principles apply when thousands of workload elements are analyzed with the program.

3.3.4.2 Scaling

One problem that is introduced by the use of statistical clustering algorithms is scaling. In the clustering example discussed in the previous section, a similar numerical range for the X and Y variables was selected. Unfortunately, actual workload data presents a much wider range of values.

Consider the problem of developing job classes. The job classes in this example are based on CPU minutes and print lines. The maximum value observed for CPU minutes would be far less than 1,000 minutes, while the likelihood of printing more than 1,000 lines is very high. If the differences in these features are squared in the geometric distance equation, the number of CPU minutes used by a job would appear to be insignificant. Therefore, the variables must be scaled prior to the clustering process so that relative differences between the features have a similar influence on the geometric distance calculation.

Neugents technology for clustering operations uses a simple unit scaling technique described next.

UNIT SCALING

Unit scaling is a simple technique for solving the problems introduced by the different range of values associated with each of the resource vector's features. Most simply stated, unit scaling maps the range of values associated with each feature into the range 0 to 1 (AGR76). Equation 3 details this mapping.

$$s_i = \frac{x_i - x_{\min}}{x_{\max} - x_{\min}} \quad (\text{Eqn 3})$$

where: s_i is the scaled value of x_i

x_i is i-th observation of the x feature

$x_{\min}$ is the minimum of all x_i

$x_{\max}$ is the maximum of all x_i

max i

A modified version of UNIT scaling is currently used within Neugents technology whereby data values are scaled internally to a range of -1 to +1.

The potentially poor representation of outliers is one of the primary concerns in the selection of features for the clustering process. Criteria for selecting features for clustering are discussed in Section 3.3.2.4.

3.3.4.3 Algorithm Implementation

The Data Clustering program is comprised of these major steps:

- 1 - Select data elements (features) from the CA MICS file based on user input.
- 2 - Execute Neugents technology training phase, which examines the input data and builds a set of clusters using its own proprietary techniques.
- 3 - Neugents technology then executes the "consulting" phase, where all observations (patterns) are consulted against the trained model and each observation or pattern is assigned to its final cluster. Outliers are marginally assigned during this process.
- 4 - Clustered data, along with model performance statistics are returned to the user.

The first step in the program is the selection of the user-specified elements from a CA MICS database file. The file and database elements (features) are specified with parameters on the Data Clustering screen. The control parameters are discussed in the sections that follow.

The second step is to perform the training phase, in a manner not unlike that done in the Workload Forecasting. A model of clusters is produced from the input data using an iterative process involving a combination of hierarchical and "k-means" methods. The actual process involved is proprietary, but permits accurate and rapid cluster creation. A Relative Position value is calculated for each observation or pattern in the manner previously described for cluster evaluation purposes, and the model is then ready for the final phase.

During the next phase, the input data is reprocessed and "consulted" against the trained model. Each observation is now assigned to a permanent cluster and any outliers are marginally assigned, meaning they are put with a cluster that best describes the vector data contents.

Lastly, the final clustered data, along with model performance statistics is made available to the Data Clustering application.

3.3.5 Component Operation

This section describes the screens and parameters you can use to control data clustering. The control parameters let you define the CA MICS file you want to use as input, the time periods you want to specify, and the variables you want to use in the clustering process. Each of the control screens is discussed in the following sections.

- 1 - Data Clustering Screen
- 2 - Extended Options Menu
- 3 - Inquiry Processing Screen
- 4 - Data Element Derivation Screen
- 5 - Alternate Input File Screen
- 6 - Save Selected Input Observations Screen
- 7 - Save Cluster Output Observations Screen
- 8 - Data Filtering Parameters Screen
- 9 - Clustering Executions Options Screen

3.3.5.1 Data Clustering Screen

Figure 3-43 illustrates the Data Clustering screen.

```

/----- Data Clustering -----\
| Command ==>                               |
|                                             |
| Composing CA MICS Inquiry: CLUCIO - Data Clustering |
|                                             |
| CA MICS file ==> -                           |
| Time-span:    DETAIL                        Cycle(s) ==> __ - __ |
| Data base ==> _                            _ |
| Allow Run-Time Override of Data Base ==> Y (Y/N) and/or Cycle(s) ==> Y (Y/N) |
|                                             |
| Analysis based on      ==> _____ |
|                                             |
| Report for each cluster ==> _____ |
|                                             |
| Data Selection:                               Allow run-time override |
| Date      ==> _____ - _____ - _____ ==> N (Y/N) |
| Hour      ==> __ - __ - __ - __ - __ - __ ==> N (Y/N) |
| Zone      ==> _____ ==> N (Y/N) |
| SYSID     ==> _____ ==> N (Y/N) |
|                                             |
| Specify extended options ==> N (Y/N) |
\-----/

```

Figure 3-43. Data Clustering Screen

The Data Clustering screen contains the following fields:

- CA MICS file: The CA MICS file that contains the data elements that you wish to include in your workload characterization analysis. If you are uncertain of the correct file name, type a question mark (?) in this field to obtain a complete selection list for your installation.
- Cycle(s): The cycle or cycles of the input file that are used in the analysis. For most installations, there is one cycle of DETAIL data for each day. You may specify a maximum of 21 cycles.
- Database: The database unit identifier(s)(up to five units) that contain the specified CA MICS file. An eight-character database unit name is displayed after you specify this identifier. If you are uncertain of which database unit the file resides in, type a question mark (?) in this field to obtain a complete selection list for your installation.
- Analysis based on: The features (data elements) that are analyzed by the clustering procedure. You can define a maximum of ten features. Cluster variables must be numeric. These data elements may either be CA MICS data elements or non-CA MICS data elements. If you specify non-CA MICS data elements, you must provide the SAS code to derive them via the exit available on the Extended Options Menu, Figure 3-48.
- Report for each cluster: The data elements that will be detailed in the Resource Consumption by Cluster Reports (see Section 3.3.3.3). At a minimum, you will generally want to specify all of the data elements specified for the Analysis based on the parameter. In addition, you can specify other numeric data elements that you might find informative in the reports. For example, clusters for a job class study might be developed using CPU time, tape mounts, and print lines as features. You might also wish to report on tape EXCPs, tape drive occupied hours, disk EXCPs, and average elapsed time of the jobs assigned to the clusters.

A maximum of ten report variable names can be specified. Report variables must be numeric. These data elements may be either CA MICS data elements or non-CA MICS data elements. If non-CA MICS data elements are specified, you must provide the SAS code to derive them via the exit available on the Extended Options Menu.

Date selection:	The selection criteria for Date, Hour, Zone, and SYSID to subset the data you are examining. You must enter date values in the ddmonyy format.
SYSID	The system identification number. If SYSID is not specified, data from all systems will be invalid in the analysis.
Specify extended options:	The specification for selecting extended options. Select this option if you wish to specify additional parameters.

3.3.5.2 Extended Options Menu

Typing a Y at the "Specify extended options" prompt on the Data Clustering screen displays the Extended Options Menu shown in Figure 3-44.

```

/----- Extended Options Menu -----\
|Command ==>|
|           |
|Inquiry Step: Data Clustering|
|           |
|Specify global data selection exit      ==> N (Y/N/R)|
|Specify non-CA MICS data element derivation exit ==> N (Y/N/R)|
|           |
|Specify alternate input file            ==> N (Y/N/R)|
|Save selected input observations        ==> N (Y/N/R)|
|Save cluster output observations       ==> N (Y/N/R)|
|           |
|Specify data filtering parameters      ==> N (Y/N/R)|
|Specify clustering execution options    ==> N (Y/N/R)|
|Define report description element names ==> N (Y/N/R)|
|           |
|           |
|           |
|           |
\-----/

```

Figure 3-44. Extended Options Menu

You can select any of the options available on this screen by typing a Y (YES) at any of the prompts. Typing an R (RESET) resets the option and nullifies any specifications you made for that option.

A description of the Extended Options unique to the Data Clustering query follow. For additional information on using the common Extended Options refer to Section 1.4.3.6, Specifying Common Extended Options.

- 1 - Save Cluster Output Observations Screen
- 2 - Data Filtering Options Screen
- 3 - Clustering Execution Options Screen
- 4 - Report Description Elements Screen

[illegible]

Data Clustering has the ability to save the SAS cluster observations from the SAS clustering program to external files for additional processing. The Save Cluster Output Observations screen allows you to specify the SAS data set and the name of the SAS file to which the cluster observations are to be written. Note that the SAS data set must already be created.

3.3.5.2.2 Data Filtering Options Screen

Figure 3-46. Data Filtering Options Screen.

```

/----- Data Filtering Options for Neugents Technology ----- ROW 1 OF 5 \
|Command ==>                                         Scroll ==> CSR |
|
|Inquiry Step:  Data Clustering                      |
|
|  Please enter the names of elements that should be checked for zero |
|  values. Observations with elements specified below that contain zero |
|  values will be subsequently discarded from further processing.      |
|
|Line Cmds: I Insert  D Delete  M Move                |
|
|Cmd Element      Data element label (long name)      Warning          |
| - -----      - - - - - - - - - - - - - - - - - - - - - - - - - - - |
| - -----      |                                     |                 |
| - -----      |                                     |                 |
| - -----      |                                     |                 |
| - -----      |                                     |                 |
| - -----      |                                     |                 |
| - -----      |                                     |                 |
| - -----      |                                     |                 |
| - -----      |                                     |                 |
| - -----      |                                     |                 |
| - -----      |                                     |                 |
\-----/

```

Figure 3-46. Data Filtering Parameters Screen

This screen allows you to specify deletion for any input observation that has a zero or missing value for a specified data element name when it is read from the CA MICS database. For example, in a job class study you might wish to delete the observations for any jobs whose TCB CPU time is zero or missing. To do so, you would code JOBTCBTM in the Element field of the display table at the bottom of the screen. You can specify only numeric variables. To obtain additional lines, enter an I in the Cmd column of any row in the display portion of the screen.

3.3.5.2.3 Clustering Execution Options Screen

Figure 3-47 illustrates the Clustering Execution Options screen.

```

/----- Execution Options for Neugents Technology ----- ROW 1 OF 5 \
|
|Inquiry Step:  Data Clustering
|
|      Clustering Input Data ==> SAMPLE      (SAMPLE/POPULATION)
|      Clustering Method   ==> RADIUS      (RADIUS/CLUSTERS)
|      Maximum Cluster Radius ==> 3.0      (1.0-3.5) If RADIUS Specified
|      Target Number of Clusters ==>        (03-64)  If CLUSTER Specified
|
|      Training Sample Size ==> 2000      (nnnnnn/ALL)
|      Sample Trim Limit   ==> 02.5      (0.00-99.9) Default: 2.5 Pct.
|
| Include Outliers in Clusters ==> YES      (YES/NO)  Trim 02.5 Pct.
| Include Sparse Clusters ==> YES      (YES/NO)  LT 0.5 Pct. of Obs.
|
|      Report Cluster Population ==> NO      (YES/NO)
| Report Using Account Fields ==> YES      (YES/NO)
| Report Outliers Separately ==> NO      (YES/NO)
|      Report Sparse Clusters ==> NO      (YES/NO)
| Create Cluster Index Graph ==> NO      (YES/NO)
| Create Population Graph ==> NO      (YES/NO)
|
|-----/

```

Figure 3-47. Clustering Executions Options Screen

This screen is used to enter the execution options for the clustering process and contains the following fields:

CA MICS file: The CA MICS file that contains the data elements that you wish to include in your data clustering analysis. If you are uncertain of the correct file name, type a question mark (?) in this field to obtain a complete selection list for your installation.

Clustering Input Data: Either SAMPLE or POPULATION, default is SAMPLE. This parameter informs the application to select either a SAMPLE of the input data, or to use the entire POPULATION for processing.

Clustering Method:	Either RADIUS or CLUSTER, default value is RADIUS. The clustering operation will be controlled by this option. If RADIUS is specified, then clustering will be limited to a certain cluster size or RADIUS. If CLUSTER is chosen, then clustering will produce only a certain number of clusters, but the clusters will vary more in size and can be quite large in scope.
Maximum	A range from 1.0 to 3.5, with a default value of 3.0. The maximum size of the largest cluster to be formed, specified in terms of standard deviations. Clustering will stop once any cluster exceeds this value. Use of this option will cause more clusters to be formed but will generally result in a better 'fit' of the data.
Target Number of Clusters:	If CLUSTER is specified for the Clustering Method, then a number from 3 to 64 must be entered here, with a default value being 8. This is the number of clusters that will be produced in a particular execution. This option is useful when using clustering as a data analysis tool where you might want to simply partition the data into heterogeneous groups, but are not concerned about the size or exact content of the clusters.
Training Sample Size:	A numeric value up to 999,999 or the word 'ALL', if ALL data is to chosen for the sample. The default value is 2,000 observations. The actual sample size is computed to be 2% of the sample or 2,000 cases, whichever is greater.
Sample Trim Limit:	A percentage of the sample to be removed, or 'trimmed' to improve the normality of the data and allow for better descriptive statistics to be obtained. The default value is 2.5% and will cause those observations that contain the top 2.5% of values for each cluster feature to be removed. This should effectively remove observations that contain very large values that tend to distort descriptive statistics, and should improve the shape and fit of the data within the clusters.

Include Outliers in Clusters:	YES/NO, with a default value of YES. This option indicates whether observations that are deemed to be outliers, that is they contain one or more feature values that exceed the trim limit. Including these observations improves the quality of the study by ensuring that all sampled data will be accounted. Outliers may be excluded if they have been reviewed and a determination is made that they are not relevant.
Include Sparse Clusters:	YES/NO, with a default value of NO. This option controls whether 'sparse' clusters are included for subsequent reporting or other post-processing operations. Sparse clusters are defined as those clusters that contain less than 0.5% of the sample population. Note, if you select the default value of NO, sparse clusters are still created and the data remains available, but the data is merely excluded from various reports and graphs.
Report Cluster Population:	YES/NO, with the default being NO. This option lets you decide if you want to generate this report. This report contains detailed usage statistics for each included cluster, with both cluster features and optional reporting elements being reported.
Report Using Account Fields:	YES/NO, with the default value being YES. Controls whether user defined CA MICS account fields are included in detail reports. Inclusion of these fields improves the content and usability of detailed reports but can result in a print line overflow and distortion of the report format.
Report Outliers Separately:	YES/NO, with the default value being NO. Controls whether to generate a detail report of all observations that contain one or more outlier cluster feature values. This report can be used to review and investigate outliers and their possible effect on the clustering operation.
Report Sparse Clusters:	YES/NO, with the default value being NO. Controls whether to generate a detail report of all 'sparse' clusters and their observations.

This report can be used to better understand the sparse clusters and to determine whether their content should be considered for additional processing.

Create YES/NO, with the default value being NO.
Cluster Controls whether to generate a printer graph of
Index the cluster index values. This graph provides
Graph: a high-level pictorial overview of cluster
performance for the study.

Create YES/NO, with the default value being NO.
Population Controls whether to generate a printer graph of
Graph: the cluster population. This graph provides a
high-level pictorial overview of cluster
population for the study.

3.3.5.2.4 Report Description Elements Screen

Figure 3-48 illustrates the Report Description Elements screen.

```

/----- Report Description Elements ----- ROW 1 OF 5\
|Command ==>                               Scroll ==> CSR |
|
|Inquiry Step:  Data Clustering                |
|
|Line cmds: I Insert  D Delete  M Move         |
|
|Cmd Element   Data element label (long name)   Warning |
| - - - - -    - - - - - - - - - - - - - - -    - - - - - |
| - - - - -    - - - - - - - - - - - - - - -    - - - - - |
| - - - - -    - - - - - - - - - - - - - - -    - - - - - |
| - - - - -    - - - - - - - - - - - - - - -    - - - - - |
| - - - - -    - - - - - - - - - - - - - - -    - - - - - |
|
\-----/

```

Figure 3-48. Report Description Elements Screen

The Report Description Elements screen allows you to specify additional descriptive data elements to be included in the optional SAS file that can be created by Data Clustering. Unlike the cluster, zero deletion, and report variables, the descriptive variables need not be numeric. For example, in a job class study you might wish to include additional data elements like JOBCLASS (job class), JOB (job name), and RDRTS (reader time stamp) for use in your SAS programs that process the optional SAS File.

3.3.5.3 Inquiry Processing Screen

Figure 3-49 illustrates the Inquiry Processing screen.

```

/----- Inquiry Processing -----\
|Command ==>|
|           |
|   Active MICF Inquiry: DCCPU - Data Clustering   |
|           Author:      M. Patriarco               |
|           |                                           |
|   Specify Inquiry processing ==> EXECUTE (EXECUTE/MODIFY/SAVE/CANCEL) |
|   EXECUTE - Execute the inquiry.                   |
|   MODIFY  - Redisplay the inquiry for modification. |
|   SAVE    - Exit from inquiry composition and save the inquiry. |
|   CANCEL  - Exit from inquiry composition without saving the inquiry. |
|           |                                           |
|Instructions:|
|           |                                           |
|   Specify desired option and press ENTER key.      |
|           |                                           |
|           |                                           |
|           |                                           |
|           |                                           |
\-----/

```

Figure 3-49. Inquiry Processing Screen

Typing END at the Command prompt or pressing the END PF key (usually PF3) while on the Data Clustering screen will display the Inquiry Processing screen. Your available options are described in the text of the screen. Note that you may save a partially composed inquiry for later modification and execution.

3.3.5.4 Data Element Derivation Screen

If you require user-defined data elements, use the Data Element Derivation Macro Definition screen shown in Figure 3-50 to specify their derivation.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> CSR |
|
|Inquiry Step: Data Clustering                    |
|
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy |
|
|Cmd  SAS Statements for: Specifying data element derivation |
| - ----- |
| _ - %MACRO DRVMAC; |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| - ----- |
| _ - %MEND DRVMAC; |
|***** BOTTOM OF DATA *****|
|
\-----/

```

Figure 3-50. Data Element Derivation Macro Definition Screen

Data Clustering provides a user-exit for defining the procedures for computing data elements that are not part of the CA MICS database. This exit is invoked by coding SAS statements in the DRVMAC macro. This macro is bracketed by the %MACRO and %MEND statements, and field protected from modification. The default for the DRVMAC macro is null.

You must code each data element derivation SAS statement in accordance with the syntax of the SAS Macro Language, as described in the SAS manuals.

```
/----- Alternate Input File -----\
|Command ==>
|
|Inquiry Step: Data Clustering
|
|SAS file name(s) ==> _____
|
|
|SAS Data set allocation instructions:
|   Data set name    ==> _____
|   Disposition      ==> SHR  (SHR/OLD)
|
|
|
|
|
|
|
|
|
|
```

Data Clustering has the ability to read data from either CA MICS files or external SAS files. The Alternate Input File screen allows you to specify that the input data for Data Clustering is a SAS file or files, and input the SAS data set name and up to five SAS files to use as input to the analysis. Note that the SAS data must already be created.

3.3.5.6 Save Selected Input Observations Screen

Figure 3-52 illustrates the Save Selected Input Observations screen.

```
/----- Save Selected Input Observations -----  
|Command ==>  
|  
|Inquiry Step: Data Clustering  
|  
|SAS file name      ==> _____  
|  
|SAS data set name   ==> _____  
|  
|  
|  
|  
|  
|  
|  
|  
|  
|  
|\-----
```

Figure 3-52. Save Selected Input Observations Screen

Data Clustering has the ability to save the observations that your analysis inquiry selected to an external SAS file for further processing by independent SAS programs, or for future processing by the Data Clustering software. The Save Selected Input Observations screen allows you to specify that the selected observations are to be written to an external SAS file, and to specify the SAS data set name and the name of the SAS file to which they are to be written. Note that the SAS data set must already be created.

[illegible]

Data Clustering has the ability to save the SAS cluster observations from the SAS clustering program to external files for additional processing. The Save Cluster Output Observations screen allows you to specify the SAS data set and the name of the SAS file to which the cluster observations are to be written. Note that the SAS data set must already be created.

3.3.5.8 Data Filtering Parameters Screen

Figure 3-54 illustrates the Data Filtering Parameters screen.

```

/----- Data Filtering Parameters for Neugents Technology ----- ROW 1 OF 5 \
|Command ==>                                                    Scroll ==> CSR |
|
|Inquiry Step:  Data Clustering                                     |
|
|  Please enter the names of elements that should be checked for zero |
|  values. Observations with elements specified below that contain zero |
|  values will be subsequently discarded from further processing.      |
|
|Line Cmds: I Insert  D Delete  M Move                             |
|
|Cmd Element      Data element label (long name)                  Warning |
| - -----      - - - - - - - - - - - - - - - - - - - - - - - - - - - |
| - -----      |
| - -----      |
| - -----      |
| - -----      |
| - -----      |
| - -----      |
| - -----      |
| - -----      |
| - -----      |
| - -----      |
|-----/

```

Figure 3-54. Data Filtering Parameters Screen

This screen allows you to specify deletion for any input observation that has a zero or missing value for a specified data element name when it is read from the CA MICS database. For example, in a job class study you might wish to delete the observations for any jobs whose TCB CPU time is zero or missing. To do so, you would code JOBTCBTM in the Element field of the display table at the bottom of the screen. You can specify only numeric variables. To obtain additional lines, enter an I in the Cmd column of any row in the display portion of the screen.

3.3.5.9 Clustering Execution Options Screen

Figure 3-55 illustrates the Clustering Execution Options screen.

```

/----- Execution Options for Neugents Technology ----- ROW 1 OF 5 \
|
|Inquiry Step:  Data Clustering
|
|      Clustering Input Data ==> SAMPLE      (SAMPLE/POPULATION)
|      Clustering Method   ==> RADIUS      (RADIUS/CLUSTERS)
|      Maximum Cluster Radius ==> 3.0      (1.0-3.5) If RADIUS Specified
|      Target Number of Clusters ==>        (03-64)  If CLUSTER Specified
|
|      Training Sample Size ==> 2000      (nnnnnn/ALL)
|      Sample Trim Limit   ==> 02.5      (0.00-99.9)  Default: 2.5 Pct.
|
| Include Outliers in Clusters ==> YES      (YES/NO)  Trim 02.5 Pct.
| Include Sparse Clusters ==> YES      (YES/NO)  LT 0.5 Pct. of Obs.
|
|      Report Cluster Population ==> NO      (YES/NO)
|      Report Using Account Fields ==> YES      (YES/NO)
|      Report Outliers Separately ==> NO      (YES/NO)
|      Report Sparse Clusters ==> NO      (YES/NO)
|      Create Cluster Index Graph ==> NO      (YES/NO)
|      Create Population Graph ==> NO      (YES/NO)
|
|-----/

```

Figure 3-55. Clustering Executions Options Screen

This screen is used to enter the execution options for the clustering process and contains the following fields:

CA MICS file: The CA MICS file that contains the data elements that you wish to include in your data clustering analysis. If you are uncertain of the correct file name, type a question mark (?) in this field to obtain a complete selection list for your installation.

Clustering Input Data: Either SAMPLE or POPULATION, default is SAMPLE. This parameter informs the application to select either a SAMPLE of the input data, or to

use the entire POPULATION for processing.

Clustering Method:	Either RADIUS or CLUSTER, default value is RADIUS. The clustering operation will be controlled by this option. If RADIUS is specified, then clustering will be limited to a certain cluster size or RADIUS. If CLUSTER is chosen, then clustering will produce only a certain number of clusters, but the clusters will vary more in size and can be quite large in scope.
Maximum	A range from 1.0 to 3.5, with a default value of 3.0. The maximum size of the largest cluster to be formed, specified in terms of standard deviations. Clustering will stop once any cluster exceeds this value. Use of this option will cause more clusters to be formed but will generally result in a better 'fit' of the data.
Target Number of Clusters:	If CLUSTER is specified for the Clustering Method, then a number from 3 to 64 must be entered here, with a default value being 8. This is the number of clusters that will be produced in a particular execution. This option is useful when using clustering as a data analysis tool where you might want to simply partition the data into heterogeneous groups, but are not concerned about the size or exact content of the clusters.
Training Sample Size:	A numeric value up to 999,999 or the word 'ALL', if ALL data is to chosen for the sample. The default value is 2,000 observations. The actual sample size is computed to be 2% of the sample or 2,000 cases, whichever is greater.
Sample Trim Limit:	A percentage of the sample to be removed, or 'trimmed' to improve the normality of the data and allow for better descriptive statistics to be obtained. The default value is 2.5% and will cause those observations that contain the top 2.5% of values for each cluster feature. This should effectively remove observations that contain very large values which tend to distort descriptive statistics and should

	improve the shape and fit of the data within the clusters.
Include Outliers in Clusters:	YES/NO, with a default value of YES. This option indicates whether observations that are deemed to be outliers, that is they contain one or more feature values that exceed the trim limit. Including these observations improves the quality of the study by ensuring that all sampled data will be accounted. Outliers may be excluded if they have been reviewed and a determination is made that they are not relevant.
Include Sparse Clusters:	YES/NO, with a default value of NO. This option controls whether 'sparse' clusters are included for subsequent reporting or other post-processing operations. Sparse clusters are defined as those clusters that contain less than 0.5% of the sample population. Note, if you select the default value of NO, sparse clusters are still created and the data remains available, but the data is merely excluded from various reports and graphs.
Report Cluster Population:	YES/NO, with the default being NO. This option lets you decide if you want to generate this report. This report contains detailed usage statistics for each included cluster, with both cluster features and optional reporting elements being reported.
Report Using Account Fields:	YES/NO, with the default value being YES. Controls whether user defined CA MICS account fields are included in detail reports. Inclusion of these fields improves the content and usability of detailed reports but can result in a print line overflow and distortion of the report format.
Report Outliers Separately:	YES/NO, with the default value being NO. Controls whether to generate a detail report of all observations that contain one or more outlier cluster feature values. This report can be used to review and investigate outliers and their possible effect on the clustering operation.
Report	YES/NO, with the default value being NO.

Sparse Clusters:	Controls whether to generate a detail report of all 'sparse' clusters and their observations. This report can be used to better understand the sparse clusters and to determine whether their content should be considered for additional processing.
Create Cluster Index Graph:	YES/NO, with the default value being NO. Controls whether to generate a printer graph of the cluster index values. This graph provides a high-level pictorial overview of cluster performance for the study.
Create Population Graph:	YES/NO, with the default value being NO. Controls whether to generate a printer graph of the cluster population. This graph provides a high-level pictorial overview of cluster population for the study.

3.3.6 Case Studies

This section provides a case study to illustrate how you can apply the data clustering methodology to solve performance management problems that many data processing installations commonly encounter.

The batch job class case study shows how you can use clustering to investigate patterns in the batch jobs being processed. You can use information to develop workload classes which can then form a reasonable basis for service objectives. The statistical uniformity of the classes enables you to assess the degree to which the service objective can be met. You can use the same information to define a job class structure for establishing an initiator scheme that is consistent with the service objectives. A full set of reports and an explanation for each report is provided with this study.

For more information, see the following subsection:

1 - Batch Job Class Study

3.3.6.1 Batch Job Class Study

The Batch Job Class case study applies the Data Clustering program to identify job classes for an installation. The following topics are addressed by section:

- 1 - Problem Description
- 2 - Control Parameters
- 3 - Report Interpretation

3.3.6.1.1 Problem Description

This problem asks you to determine a new set of resource-limited job classes for an installation. Job class structures are generally established to meet two objectives (FER78, ART79a):

- o Provide operational controls
- o Facilitate scheduling and implement service policies

You usually establish job classes that you define to meet the first objective for reasons other than resource requirements. For example, online systems like IMS and CICS are often run in unique job classes to avoid interference with other workloads. Also, many computer centers use one or two job classes to segregate production jobs from the remainder of their workload.

Job classes established to meet the second objective provide users with alternative classes for job assignments based on resource requirements. Typically, the service objective of a job class decreases (that is, jobs have longer turnaround times) as its resource requirements increase. Using a job class structure that focuses on scheduling and service has a number of potential benefits:

- o Many jobs can be completed in a short time (assuming demand) since the job classes representing smaller jobs may be expedited.
- o The availability of better turnaround time for the smaller classes may encourage users to conserve resources.
- o Job scheduling is simplified since each job class represents a homogeneous set of jobs.

Clustering is an ideal approach to this problem since we want to exploit the natural characteristics of the workload. If we arbitrarily select job classes, we might select classes that split normal resource consumption patterns or choose some that are too large or small to be useful.

For this sample installation, we investigated job classes based on CPU time and tape drives usage. We selected these resources to quantify the system, output, and operator requirements presented by the jobs. Additionally, several other reporting metrics were included to quantify the resources represented by each cluster. Although a number of other features could have been selected for analysis, only

these were chosen to keep the resulting job class structure as simple as possible.

3.3.6.1.2 Control Parameters

Figure 3-50 illustrates the Data Clustering screen.

```

/----- Data Clustering -----\
|Command ==>
|
|Composing CA MICS Inquiry: CLASS - Batch Job Class Case Study
|
|CA MICS file ==> JOB -
|Time-span:      DETAIL                      Cycle(s) ==> 01 - 10|
|Data base ==> A
|Allow Run-time Override of Data Base ==> Y (Y/N) and/or Cycle(s) ==> Y (Y/N)|
|
|Analysis based on      ==> JOBTCBTM  TAPEUSE_  _____  _____  _____|
|
|Report for each cluster ==> JOBEDASD  JOBETAPE  JOBTOHRS  JOBNLR  _____|
|
|_____  _____  _____  _____  _____|
|
|Data Selection:                      Allow run-time override|
|Date      ==> _____ - _____ - _____ ==> N (Y/N)|
|Hour      ==> __ - __ - __ - __ - __ ==> N (Y/N)|
|Zone      ==> 1 2 3 4 _ _ _ _ _ ==> N (Y/N)|
|SYSID     ==> SYSA _____ ==> N (Y/N)|
|
|Specify extended options ==> Y (Y/N)
\-----/

```

Figure 3-50. Data Clustering Screen

The Data Clustering screen contains the following fields:

CA MICS file: The CA MICS file to be used for workload characterization. (For this example, we specified the CA MICS BATJOB File for a job class study.)

Cycles: The number of cycles to be specified. (For example, we specified ten cycles. In an actual study, analyze as much data as possible before you initiate a new job class structure.)

Database: The single-character database identifier for this file at your site.

Analysis based on: The CA MICS data element feature names (CPU Time and Tape Drive Usage) to be used in the clustering procedure. Because TAPEUSE is a user supplied element that is used to summarize tape device usage, it is defined on the User-Defined Data Element Derivation Macro screen and uses the DRVMAC macro to populate the data values. (As discussed in Section 3.3.2.4, you should specify as few features as possible.)

DASD EXCPs, Tape EXCPs, Tape Drive Occupancy Time, and print lines are specified for reporting purposes.

```
/----- Extended Options Menu -----\
|Command ==>
|
|Inquiry Step: Data Clustering
|
|
|Specify global data selection exit          ==> N (Y/N/R)
|Specify non-CA MICS data element derivation exit==> Y (Y/N/R)
|
|Specify alternate input file                ==> N (Y/N/R)
|Save selected input observations            ==> N (Y/N/R)
|Save cluster output observations           ==> Y (Y/N/R)
|
|Specify data filtering parameters           ==> Y (Y/N/R)
|Specify clustering execution options        ==> Y (Y/N/R)
|Define report description element names     ==> N (Y/N/R)
|
```



Figure 3-51. Data Clustering Extended Options Menu

You can select additional option specifications on the Extended Options Menu. Figure 3-51 shows an example with four of the extended options specified.




Figure 3-52. Report Description Elements Screen

Since we wished to create an optional external report data set during the study, we also specified two descriptive variables. We selected the variable, JOBCCLASS (job class), so that the results of the cluster analysis could be compared with the current job class structure and the assignment of particular jobs could be determined. These data elements are specified on the Report Description Elements screen shown in Figure 3-52 above.

```

/----- SAS Statements -----\
|Command ==>                               Scroll ==> CSR |
|
|Inquiry Step:  Data Clustering
|
|Line Cmds: I Insert  D Delete  R Repeat  M Move  C Copy
|
|Cmd  SAS Statements for:  Specifying data element derivation
|  -  -----
|  - %MACRO DRVMAC;
|    IF JOBMXNTA EQ 0 THEN TAPEUSE = 0;
|    ELSE IF 01 LE JOBMXNTA LE 3 THEN TAPEUSE = 1;
|    ELSE TAPEUSE = 2;
|    LABEL TAPEUSE = "Tape Device Usage Indicator";
|  - %MEND DRVMAC;
|***** BOTTOM OF DATA *****
|
|
|
\-----/

```

Figure 3-53. User-Defined Data Element Derivation Macro

The SAS statements that compute the user-defined data elements are placed in the DRVMAC macro shown in Figure 3-53 above.

```

/----- Data Filtering Options ----- ROW 1 OF 1\
|Command ==>                               Scroll ==> CSR |
|
|Inquiry Step:  Data Clustering                |
|
|  Please enter the names of elements that should be checked for zero
|  values. Observations with elements specified below that contain zero
|  values will be subsequently discarded from further processing.
|
|Line Cmds:%I+Insert %D+Delete %M+Move
|
|Cmd Element      Data element label (long name)      Warning
| - - - - -      - - - - -
|  _  JOBTCBTM
|  _
|  _
|***** BOTTOM OF DATA *****
|
|
\-----/

```

Figure 3-54. Filter Observations

In this case study we wanted to create the SAS data set name, ALLOBS, that contains the cluster analysis output. To do this, we specified ALLOBS as the SAS file name on the External Report Output File Definition screen shown in Figure 3-55 above.

```

----- Inquiry Processing -----
|Command ===>
|
|
|Active MICF Inquiry: CLASS - Batch Job Class Case Study
|                Author: Charles Frost
|
|Specify Inquiry processing ===> EXECUTE (EXECUTE/MODIFY/SAVE/CANCEL)
|EXECUTE - Execute the inquiry.
|MODIFY - Redisplay the inquiry for modification.
|SAVE - Exit from inquiry composition and save the inquiry.
|CANCEL - Exit from inquiry composition without saving the inquiry.
|
|Instructions:
|
|Specify desired option and press ENTER key.

```

Figure 3-56. Inquiry Processing

After a Data Clustering analysis inquiry is created, the Inquiry Processing screen shown here is displayed. From this screen, you can execute the newly created inquiry, modify it, save it without executing it, or cancel the parameter specifications you just entered.

3.3.6.1.3 Report Interpretation

The example we discussed in Section 3.3.6.1.2 showed you how to use the control screens to execute Data Clustering. The program produces five reports and two graphs, of which four will be used in this study.

- 1 - Descriptive Statistics Report
- 2 - Cluster Performance Summary Report
- 3 - Cluster Population Summary Report
- 4 - Data Value Exception Detail Report
- 5 - Analysis of the Clustering Output Data Set

Additionally, a separate section will explain how you can further analyze the output data produced during clustering and explore possible job class scenarios derived from the case study.

3.3.6.1.3.1 Descriptive Statistics Report**3.3.6.1.3.2 Clustering Performance Summary Report****3.3.6.1.3.3 Clustering Population Summary Report****3.3.6.1.3.4 Data Value Exception Detail Report**

This report provides a simple listing of all observations that are considered outliers, including a flag indicating which cluster feature(s) value caused the outlier status. This report can be used to help verify the criteria used for clustering, to identify those cases that should be considered as mainstream, despite the outlier status, and to determine how such cases should be considered in any post-processing efforts.

Data Clustering Analysis
Data Value Exception Detail
For: Tuesday, July 1, 2003

1

Listing of Observations With Outlier Values

System Identifier	DIVISION	DEPARTMENT	RESERVED	Job Group	Job Identification	End Time Stamp	Job TCB CPU Time	Tape Device Usage Indicator
XAL1	DEV	01056	*****	7	LTSIPCSX	14NOV02:17:54:24.33	0:40:00.28*	0
		01063	*****	5	HEMDA023	09NOV02:03:50:34.20	0:06:00.34*	0
	GIS	00045	*****	0	A02DDSTT	07NOV02:21:15:40.57	4:56:12.13*	0
						09NOV02:20:05:43.32	0:32:20.13*	0
						12NOV02:21:58:37.99	1:24:08.19*	0
					A02ICSTT	07NOV02:21:15:41.47	0:14:03.47*	0
					A02SGATE	09NOV02:20:03:54.73	0:11:27.99*	0
					DISKIB03	12NOV02:02:52:05.00	0:17:34.86*	1
						14NOV02:23:43:56.22	0:08:37.51*	1
					DISKIB04	14NOV02:22:42:28.48	0:07:55.69*	1
					DISKIB08	05NOV02:07:26:09.92	0:25:25.81*	1
					MICSDDAP	04NOV02:04:17:41.38	0:22:59.33*	1
					MICSDDAV	04NOV02:02:18:04.06	0:03:46.72	2*
					TOPCDEV2	08NOV02:09:51:53.06	0:05:40.66*	0
					TOPDDEVL	06NOV02:17:20:43.25	0:28:48.32*	0
						09NOV02:20:06:02.62	0:12:11.98*	0
					TOPDDEV2	12NOV02:21:58:38.00	0:53:26.01*	0
						13NOV02:10:25:47.61	0:20:52.34*	0
					TOPDQAQA	06NOV02:11:35:54.43	0:17:50.60*	0
						12NOV02:21:58:37.99	0:07:03.85*	0
					TOPSPCDB	05NOV02:20:03:27.63	4:20:09.74*	0
						09NOV02:20:04:14.55	1:12:34.37*	0
						10NOV02:20:03:10.09	4:13:26.19*	0
						14NOV02:20:03:18.98	5:01:29.65*	0
					WKBKU2	09NOV02:16:32:04.49	0:16:44.47*	1
					WKBKU24	09NOV02:22:42:12.68	0:20:37.46*	1
					WKBKU27	10NOV02:01:31:32.87	0:24:41.93*	1
					WKBKU3	10NOV02:05:39:32.77	0:11:02.99*	1
					WKBKU30	10NOV02:06:14:29.78	0:21:42.96*	1
					WKBKU40	10NOV02:08:37:18.76	0:20:39.41*	1
					WKBKU61	10NOV02:14:51:03.13	0:16:30.64*	1
					WKBKU63	10NOV02:16:19:18.56	0:16:41.95*	1
					WKBKU67	10NOV02:19:43:18.32	0:19:06.31*	1
					WKBKU69	10NOV02:22:06:54.12	0:17:34.56*	1
					WKBKU80	09NOV02:08:03:40.48	0:13:40.25*	1
					WKBKU81	11NOV02:00:10:54.65	0:13:22.97*	1
					WKBKU86	04NOV02:02:23:29.11	0:15:29.06*	1
						11NOV02:04:51:03.00	0:17:21.82*	1
					WKBKVM81	04NOV02:14:41:56.00	0:19:28.42*	1
						11NOV02:08:57:18.00	0:19:25.72*	1
					WKBKVM84	11NOV02:10:52:19.00	0:20:15.46*	1
					WKBKVM85	05NOV02:03:42:10.00	0:25:11.81*	1
						11NOV02:12:50:10.00	0:25:27.51*	1
					WKBKVM86	05NOV02:04:04:08.00	0:23:29.65*	1
		02046	*****	0	T1PR2NEW	14NOV02:12:56:45.00	0:07:12.03*	0
				7	MELRI01A	09NOV02:18:24:56.00	0:20:52.65*	0
					MELRI01B	07NOV02:18:12:06.00	0:25:34.91*	0
	UNK	UNK	*****	0	CACTAPG	07NOV02:21:17:12.00	0:18:03.35*	0

Note: '*' Indicates that the data value is an outlier

Data Clustering Analysis
Data Value Exception Detail
For: Tuesday, July 1, 2003

2

Listing of Observations With Outlier Values

System	Job	Job	Job TCB CPU	Tape Device Usage

Identifier	DIVISION	DEPARTMENT	RESERVED	Group	Identification	End Time Stamp	Time	Indicator
XAL1	UNK	UNK	*****	0	CACTUSA	12NOV02:21:58:38.00	0:06:32.16*	0
					TNGWE2	09NOV02:20:08:39.48	0:06:52.27*	0
					TOPDEM	12NOV02:21:58:38.00	0:25:24.12*	0

Figure 3-61. Cluster Performance Index graph

Note that in all but one case, the observations were outliers because of abnormal TCB Time usage. Such jobs can be handled either by a special exception job class, or by establishing WLM Service class definitions that manage such jobs in a lower priority scheme, such as making them Discretionary.

3.3.6.1.3.5 Analysis of the Clustering Output

The final step of the job class identification process involves some decisions based on the cluster descriptions and the analyst's experience. Although the clusters developed in the example provide a reasonable representation of the system's workload, additional work is needed to define an appropriate job class structure.

After reviewing the Cluster Performance and Population Summary reports, a simple analysis involving selecting, sorting and printing was executed against the output statistical data from the clustering operation. The output of the analysis is presented next.

Listing of Selected Data Clustering Cases Presented by Clustering Feature Values												1
Cluster Type	Minimum TAPEUSE Value	Mean TAPEUSE Value	Maximum TAPEUSE Value	Minimum JOBTCBTM Value	Mean JOBTCBTM Value	Maximum JOBTCBTM Value	Cluster Number	Cluster Radius	Total Observations	Normal Observations	Total Outlier Observations	Total Observations
Normal	0	0.00000	0	0:00:00.01	0:00:03.13	0:00:41.26	9	.190520	1,840		0	1,840
	0	0.00000	0	0:00:46.15	0:01:23.66	0:02:09.66	11	.436100	31		0	31
	0	0.00000	0	0:03:04.29	0:04:40.17	0:05:32.62	18	.993240	7		0	7
	1	1.00000	1	0:00:00.03	0:00:19.22	0:01:23.85	14	.017900	46		0	46
	1	1.00000	1	0:01:45.20	0:02:40.05	0:03:00.43	13	.712520	12		0	12
	1	1.00000	1	0:03:31.30	0:03:38.57	0:03:50.19	5	.362838	5		0	5
	1	1.00000	1	0:04:15.47	0:04:49.99	0:05:28.36	10	.197910	6		0	6
									1,947		0	1,947
Listing of Selected Data Clustering Cases Presented by Clustering Feature Values												2
Cluster Type	Minimum TAPEUSE Value	Mean TAPEUSE Value	Maximum TAPEUSE Value	Minimum JOBTCBTM Value	Mean JOBTCBTM Value	Maximum JOBTCBTM Value	Cluster Number	Cluster Radius	Total Observations	Normal Observations	Total Outlier Observations	Total Observations
Other	0	0.00000	0	0:05:40.33	0:06:25.95	0:07:12.03	12	.438660	1		6	7
	0	0.00000	0	0:11:27.99	0:11:49.98	0:12:11.98	7	.686676	0		2	2
	0	0.00000	0	0:17:50.60	0:17:56.97	0:18:03.35	4	.199029	0		2	2
	0	0.00000	0	0:20:52.34	0:20:52.49	0:20:52.65	2	.004837	0		2	2
	0	0.00000	0	0:25:24.12	0:25:29.51	0:25:34.91	3	.168434	0		2	2
	1	1.00000	1	0:07:55.69	0:08:16.60	0:08:37.51	6	.652809	0		2	2
	1	1.00000	1	0:13:22.97	0:14:10.76	0:15:29.06	15	.444520	0		3	3
	1	1.00000	1	0:16:30.64	0:17:04.72	0:17:34.86	8	.063880	0		6	6
	1	1.00000	1	0:19:06.31	0:20:10.82	0:21:42.96	16	.876600	0		7	7
	1	1.00000	1	0:22:59.33	0:24:32.67	0:25:27.51	17	.914180	0		6	6
	0	0.21429	2	0:02:44.98	1:43:52.37	5:01:29.65	1	.000000	1		13	14
									2		51	53
Other									2		51	53

Figure 3-62. Cluster Feature Summary Statistics

The code used for this analysis divided clusters into two groups: one containing most of the normal data and the other representing the outliers. This division was done mainly to demonstrate that the outliers will normally be considered as exceptions to the job classification process. Usually only the regular (non-sparse) clusters will be processed and the sparse clusters and outliers will be treated as exceptions. In our case study, three of the sparse clusters, clusters 5, 10, and 18 were included in the "normal" clusters because they contained all normal observations. One regular cluster, cluster 1 was placed in the "Other" category because it contained almost all outliers. These decisions require installation knowledge and careful planning before implementation.

A SAS program was written to evaluate the proposed classes. The program utilizes the clustering criteria from the above report and attempts to assign each job (case) to the proposed classes (tested in ascending order of size) until a match is found. The results from the execution of this program are shown in Figure 3-63.

```
DATA JOBCLASS;
SET TOTAL_SAMPLE_CLUSTERS;

TITLE;
FOOTNOTE;

SELECT (TAPEUSE);
  WHEN (0) DO;
    IF JOBTCBTM LE 40 THEN CLASS = "A";
    ELSE IF JOBTCBTM LE 300 THEN CLASS = "B";
    ELSE CLASS = "X";
    OUTPUT;
  END;

  WHEN (1) DO;
    IF JOBTCBTM LE 300 THEN CLASS = "C";
    ELSE CLASS = "X";
    OUTPUT;
  END;

  OTHERWISE DO;
    CLASS = "X";
    OUTPUT;
  END;

END;
```

RUN;

```
PROC FREQ DATA=JOBCLASS;
```

```
  TABLES CLASS;
```

```
  FORMAT CLASS $CLASS.;
```

```
  TITLE1 "Analysis of Proposed Job Classes";
```

RUN;

The output from the job class structure evaluation is presented next.

Analysis of Proposed Job Classes

The FREQ Procedure

Proposed Job Class

Proposed Job Class	Frequency	Percent	Cumulative Frequency	Cumulative Percent

: 0 Tapes; LE 40 secs. TCB Time	1837	91.85	1837	91.85
: 0 Tapes; LE 300 secs. TCB Time	39	1.95	1876	93.80
:1-3 Tapes; LE 300 secs. TCB Time	66	3.30	1942	97.10
:<===== All Other Jobs =====>	58	2.90	2000	100.00

Figure 3-63. Analysis of Proposed Job Classes

The code above can be changed to suite the installation service and resource requirements. The proposed job class structure does a reasonably good job representing the resource requirements for the jobs being studied. The exception class ("X") contains only 58 jobs and represents less than 3.0 percent of the total number of jobs being analyzed. Changing the classification criteria in the code above might place more jobs in the exception class. This is not necessarily bad, but keep in mind that in most installations such jobs often require manual scheduling or similar support. Obviously the criteria used was selected mainly for illustrative purposes and might not be realistic given an installation's requirements. You can use the reports and substitute a variety of criteria, execute the simulation code and evaluate the results.

3.3.7 Requirements

The Data Clustering analysis tool exploits features that utilize CA's U.S. Patented Neugents(R) technology.

The exploitation of Neugents technology 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.

Appendix A: MESSAGES

Messages

This appendix lists all messages generated by the CA MICS Performance Manager Option. Most messages are caused by various conditions in the data. The messages are listed in ascending numerical sequence and include the full text of the message, the type, the reason for the message, the action you should take, and applicable document references.

Four type codes are used to categorize the messages:

1. Information: A note that documents a CA MICS Performance Manager Option or potentially important feature in the data.
2. Warning: A condition of either the data or the inquiry options that does not affect the operation of the CA MICS Performance Manager Option, but that can lead to unexpected results.
3. Error: An indication that a problem has been encountered with an inquiry option that will prevent a successful run of the product. Execution is stopped after all inquiry options have been processed.
4. Termination: An indication that a critical problem has been encountered with the data that requires the product to terminate processing immediately.

The message texts often include references to information contained in an inquiry option or values of variables. In these cases, the message text contains a string in double quotes indicating that a value will be substituted into the text at execution time.

```
+-----+
| P E R 0 0 1 2 1 |
+-----+
```

TEXT: CPU SERVICE FACTOR IS NOT DEFINED FOR:
PRSMCEC: "prsmcec"
SYSPLEX: "sysplex"
SYSID: "sysid"
CPU MODEL: "cpu model"
NUMBER OF PROCESSORS: "number"
PLEASE UPDATE THE MSU CONVERSION TABLE.

TYPE: Warning

REASON: The CPU service factor for the CPU model from
which the HARCPU data originated is not
available.

ACTION: None required.

REFERENCES: Section 2.1.5, Requirements

```
+-----+
| P E R 0 1 1 4 9 |
+-----+
```

TEXT: "number" OBSERVATIONS WERE DELETED DUE TO
MISSING VALUES.

TYPE: Information

REASON: The identified number of observations were
deleted from the analysis due to the
specifications from the control panel.

ACTION: None required.

REFERENCES: Section 3.3, Data Clustering Analysis

```
+-----+
| P E R 0 1 1 5 0 |
+-----+
```

TEXT: NO OBSERVATIONS MET SELECTION CRITERIA.
EXECUTION TERMINATES.

TYPE: Termination

REASON: The selection criteria defined via the control panel resulted in no observations being selected for analysis.

ACTION: Examine the selection criteria specified on the control panel, as well as the input data. Modify the selection criteria or select from a different CA MICS data source and resubmit.

REFERENCES: Section 3.3, Data Clustering Analysis

```
+-----+  
| P E R 0 1 1 5 1 |  
+-----+
```

TEXT: NUMBER OF CLUSTERS GENERATED EXCEEDS 99 - USE FEWER VARIABLES FOR CLUSTERING.

TYPE: Error

REASON: The variables specified for the analysis based on options on the inquiry panel resulted in more than the maximum of 99 clusters being generated.

ACTION: Specify a fewer number of variables and resubmit.

REFERENCES: Section 3.3, Data Clustering Analysis

```
+-----+  
| P E R 0 2 0 0 1 |  
+-----+
```

TEXT: NO DATA HAS PASSED THROUGH INITIAL SELECTION. CHECK FILE AND SELECTION CRITERIA VALIDITY. RUN TERMINATES WITH A USER 998 CONDITION CODE.

TYPE: Error

REASON: All data from the requested input files was either rejected by selection specifications, or the required files contained no observations to select.

ACTION: Run terminates with a user 998 abend. Verify the contents of the required files, alter selection, input specifications, or both as required, and

resubmit.

REFERENCES: None

```
+-----+  
| P E R 0 2 0 0 2 |  
+-----+
```

TEXT: THE WEEKS LEVEL HARDVA FILE IS NOT SEQUENCED BY HOUR. HARDVA RELATED MEASUREMENTS ARE UNAVAILABLE. RESPONSE TIME CANNOT BE REPORTED ON.

TYPE: Information

REASON: ASM inquiries at the WEEKS timespan, with by-the-hour reporting specified, cannot obtain HARDVA information by the hour, because the HARDVA File in the WEEKS timespan is not sequenced by hour.

ACTION: The report-by-hour is still produced, but all report elements obtained from the HARDVA File are set to missing values. You can run the report without the by-hour option to generate a report with no missing values.

REFERENCES: Section 2.3, Auxiliary Storage Management Analysis

```
+-----+  
| P E R 0 2 0 0 3 |  
+-----+
```

TEXT: MORE THAN NINE DOMAINS FOUND IN INPUT. FIRST NINE DOMAINS WILL BE PLOTTED.

TYPE: Warning

REASON: The inquiry routines allow up to nine domains to be plotted on a single page, but more than nine domains have been found in the input data. Only the first nine domains appear.

ACTION: Use the selection options on the inquiry panel to limit the number of selected domains to nine.

REFERENCES: None

+-----+
| P E R 0 2 0 0 4 |
+-----+

TEXT: NO VALID TRACE DATA FOUND, PLOTS CANNOT BE
PRODUCED. REPORT ELEMENTS DERIVED FROM TRACE
DATA ARE UNAVAILABLE.

TYPE: Warning

REASON: No RMF Trace data (CA MICS SCPTRC File) was
found. Without this data, significant portions
of the IPS output reports will be missing.

ACTION: Activate the CA MICS collection of RMF Trace
data and resubmit the inquiry. Contact your
CA MICS system administrator for assistance.

REFERENCES: None

+-----+
| P E R 0 2 0 0 5 |
+-----+

TEXT: MORE THAN 5 CLASSES WERE FOUND FOR ONE OR MORE
INITIATORS - THE REDUCED BACKLOG MODEL WILL BE
USED.

TYPE: Information

REASON: More than five job classes were specified for at
least one of the batch initiators. Because of
software restrictions, a reduced backlog model
is used when this condition is found.

ACTION: If the reduced backlog model is not acceptable,
change the inquiry so that no initiator
specifies more than five job classes.

REFERENCES: Section 3.1.6, Component Operation

+-----+
| P E R 0 2 0 0 6 |
+-----+

TEXT: JOBPRTY HAS BEEN FOUND MISSING. IT WILL BE SET
TO ZERO.

TYPE: Warning

REASON: At least one of the observations selected for processing had no valid JOBPRTY (Job Priority). The routines continue processing, but assign a JOBPRTY of zero to this job.

ACTION: None required.

REFERENCES: Section 3.1.6, Component Operation

```
+-----+  
| P E R 0 2 0 0 7 |  
+-----+
```

TEXT: NO JOBS SELECTED. SIMULATION WILL NOT BE DONE.

TYPE: Information

REASON: The simulation software cannot continue because no observations were selected for processing.

ACTION: Examine the selection criteria specified on the inquiry panel to make sure it is correct. Make sure the specified input source contains valid data. Resubmit the inquiry.

REFERENCES: Section 3.1.6, Component Operation

```
+-----+  
| P E R 0 2 0 0 8 |  
+-----+
```

TEXT: THE CLOCK STEP HAS BEEN SET TO "value" SECONDS.

TYPE: Information

REASON: The inquiry routines use a clock step value equal to that shown in the message.

ACTION: If the value in the message is not acceptable, you can specify an alternate value by using the Extended Options Menu panel of the inquiry.

REFERENCES: Section 3.1.6, Component Operation

```
+-----+  
| P E R 0 2 0 0 9 |  
+-----+
```

TEXT: JOBQ SIZE EXCEEDED - JOB INPUT STOPPED AT
"value".

TYPE: Termination

REASON: More jobs were read into the simulation than
could be handled by the default internal array
sizes. In the message text, "value" represents
the number of jobs read before the error was
detected.

ACTION: Increase the number of classes specified for one
of the initiators to at least six. If the
problem persists, contact CA Technical Support
at <http://ca.com/support>.

REFERENCES: Section 3.1.6, Component Operation

```
+-----+  
| P E R 0 2 0 1 0 |  
+-----+
```

TEXT: RUNAWAY LOOP DETECTED IN JOB-SELECTION MODULE.

TYPE: Termination

REASON: A logic error resulted in a program loop in the
job selection portion of the simulation routine.

ACTION: Contact CA Technical Support at
<http://ca.com/support>.

REFERENCES: Section 3.1.6, Component Operation

```
+-----+  
| P E R 0 2 0 1 1 |  
+-----+
```

TEXT: JOB "job", STEP NUMBER "step", PROGRAM "program"
ALLOCATED BOTH 3480 AND NON-3480 DRIVES. MOUNT
TIME ADJUSTMENT IS GLOBAL IN NATURE. ADJUSTMENT
WILL OCCUR.

TYPE: Warning

REASON: The inquiry routines determined that the job
specified in the error message used both 3420
and 3380 tape drives. The adjustment for mount
time is not as accurate as if only one type of

device had been used.

ACTION: None required.

REFERENCES: Section 2.4.7, Tape Drive Simulation Analysis Inquiry

```
+-----+
| P E R 0 2 0 1 2 |
+-----+
```

TEXT: AT LEAST ONE TYPE 30 INTERVAL RECORD HOLDING THE TAPE DRIVE ALLOCATION TOTALS WAS DROPPED.

TYPE: Information

REASON: A CA MICS observation representing an SMF Type 30 interval record showed no tape drives were allocated (PGMMXNTA=0) but tape EXCPs were recorded (PGMETAPE not 0). The invalid record is dropped from the analysis.

ACTION: None. If you get an excessive number of these messages, you may want to investigate the reason for the invalid data.

REFERENCES: Section 2.4.7, Tape Drive Simulation Analysis Inquiry

```
+-----+
| P E R 0 2 0 1 3 |
+-----+
```

TEXT: CONNECT TIME CANNOT BE ADJUSTED IN MORE THAN ONE FASHION IN THE SAME STEP RECORD. NO ADJUSTMENT WILL OCCUR FOR JOB "job" STEP "step".

TYPE: Information

REASON: Only one type of connect time adjustment is possible in one step records. Adjustment is bypassed for the indicated job and step due to a multiple adjustment attempt.

ACTION: Examine your inquiry options to make sure only one adjustment is occurring per step.

REFERENCES: Section 2.4.7, Tape Drive Simulation Analysis Inquiry

```
+-----+  
| P E R 0 2 0 1 4 |  
+-----+
```

TEXT: RECORDS WITH MISSING START TIMESTAMPS WERE
DROPPED FROM THE ANALYSIS, BAD RECORDS =
"count".

TYPE: Information

REASON: The number of observations shown in the message
were dropped from the simulation because they
did not contain a valid start date and time.

ACTION: None. If you get an excessive number of these
messages, you may want to investigate the reason
for the invalid data.

REFERENCES: Section 2.4.7, Tape Drive Simulation Analysis
Inquiry

```
+-----+  
| P E R 0 2 0 1 5 |  
+-----+
```

TEXT: RECORDS WITH MISSING END TIMESTAMPS WERE DROPPED
FROM THE ANALYSIS, BAD RECORDS = "count".

TYPE: Information

REASON: The number of observations shown in the message
were dropped from the simulation because they
did not contain a valid end date and time.

ACTION: None. If you get an excessive number of these
messages, you may want to investigate the reason
for the invalid data.

REFERENCES: Section 2.4.7, Tape Drive Simulation Analysis
Inquiry

```
+-----+  
| P E R 0 2 0 1 6 |  
+-----+
```

TEXT: THE VOLSER - "volume" IS THE ONLY VOLUME IN DAY
- "day" AT "timespan". CONTENTION INDICES
CANNOT BE COMPUTED FOR THIS TIME FRAME, RECORD

DROPPED FROM ANALYSIS.

TYPE: Information

REASON: Only one volume was selected from the day and timespan shown in the message. A contention index cannot be computed using just one observation. The observation is dropped.

ACTION: None required.

REFERENCES: Section 2.4.6, DASD Multiple Volume Analysis Inquiry

```
+-----+
| P E R 0 2 0 1 7 |
+-----+
```

TEXT: MULTIPLE VOLUME ANALYSIS ASSUMES GROUPS OF "number" DEVICES WILL BE ONLINE DURING THE SAME TIME FRAME. FOR THE LAST "value" GROUPS THIS WAS NOT THE CASE, THESE GROUPS ARE DROPPED FROM THE REPORT.

TYPE: Information

REASON: The number of online volumes found during a reporting interval was not an even multiple of the number specified for this n-way analysis. Some devices are dropped from the analysis.

ACTION: Either ignore the message or change the volumes specified in the inquiry to avoid the problem.

REFERENCES: Section 2.4.6, DASD Multiple Volume Analysis Inquiry

```
+-----+
| P E R 0 2 0 1 8 |
+-----+
```

TEXT: MULTIPLE VOLUME ANALYSIS PLOTS ARE DONE WHEN CONTENTION INDEX GT 5 OR SUMMATION INDEX GT 25. GROUP "number" DID NOT QUALIFY FOR PLOTTING.

TYPE: Information

REASON: To reduce the output produced, a plot for a volume group is produced only if the group has a

contention index greater than 5 or a summation index greater than 25. A plot was not produced for the group indicated in the message.

ACTION: None required.

REFERENCES: Section 2.4.6, DASD Multiple Volume Analysis Inquiry

```
+-----+  
| P E R 0 2 0 1 9 |  
+-----+
```

TEXT: NO SUMMARIZED OBSERVATIONS HAVE PASSED THE SPECIFIED FILTERING PROCESS. REPORT CANNOT BE PRODUCED. TEST WAS "statement".

TYPE: Error

REASON: You specified a usage qualification to apply against the observations in the inquiry that caused all observations to be deleted from the inquiry. No output can be produced.

ACTION: Remove or modify the usage qualification specified on the inquiry panel so that some data is selected for reporting.

REFERENCES: Section 2.4 - I/O Configuration Analysis

```
+-----+  
| P E R 0 2 0 2 0 |  
+-----+
```

TEXT: REPORT CANNOT BE RUN. A NON-XA RMF LEVEL HAS BEEN FOUND AND THE "file(s) are/is" NOT AVAILABLE.

TYPE: Error

REASON: Because the input data being processed is from a non-XA system, certain additional input files are needed to produce the requested inquiry. The message indicates that one or more of these files cannot be found.

ACTION: Use a different source of input data or determine why the listed files are not available.

REFERENCES: Section 2.4.5, DASD Skew Analysis Inquiry

```
+-----+  
| P E R 0 2 0 2 1 |  
+-----+
```

TEXT: THE LINES PER PAGE WAS SET TO LESS THAN 30, THIS REPORT REQUIRES A MINIMUM PAGE SIZE OF 30, PAGE SIZE SET TO 30.

TYPE: Information

REASON: Your MICF options specify a page size of less than 30 lines per page, but this report requires at least 30 lines per page for acceptable output. The page size has been set to 30 lines for this report.

ACTION: Change your MICF options to use a page size of at least 30.

REFERENCES: None

```
+-----+  
| P E R 0 2 0 2 2 |  
+-----+
```

TEXT: WARNING THE EFFECTS OF SHARED DASD MAY NOT BE CORRECTLY EVALUATED WHEN SYSID SELECTION IS IN EFFECT.

TYPE: Information

REASON: To be effective, a shared DASD report must consider all systems that share a DASD device. The inquiry being run specified that input data be selected by SYSID. This could cause the report output to be inaccurate if not all systems are considered.

ACTION: Evaluate the reasons why some systems are being excluded. Change the inquiry options if necessary.

REFERENCES: Section 2.4.3, Shared DASD Analysis Inquiry

```
+-----+  
| P E R 0 2 0 2 3 |  
+-----+
```


+-----+

TEXT: NO PR/SM LPAR DATA FOUND IN HARLPC FILE(S).
REPORT CANNOT BE PRODUCED UNTIL PR/SM DATA IS
PROCESSED BY CA MICS. RUN TERMINATES WITH A USER
998 CONDITION CODE.

TYPE: Error

REASON: An attempt was made to produce a report for
PR/SM LPAR analysis, but no PR/SM data has yet
been processed into the CA MICS database by the
standard update jobs.

ACTION: Run terminates with a user 998 abend. Check with
your system administrator as to when PR/SM data
will be available.

REFERENCES: Section 2.6, PR/SM LPAR Analysis

+-----+

| P E R 0 2 0 2 4 |

+-----+

TEXT: ALLOCATING MMG WORK DATA SET. DDNAME: MMGWRKLD/
DATA SET: "data set name"

TYPE: Information

REASON: An MMG workload data set was specified for a
graphic analysis inquiry. The system
dynamically allocates the data set and parses
the table in order to group together homogeneous
workload types. If subsequent error messages
indicate some sort of allocation problem, this
message shows the data set name that was being
allocated.

ACTION: Unless a subsequent error message indicates an
allocation problem, this message can be ignored.

REFERENCES: Section 2.7.6.3, Workload Groupings

+-----+

| P E R 0 2 0 2 5 |

+-----+

TEXT: DYNAMIC ALLOCATION ERROR FOR DATA SET: "data set
name" / ERROR CODE: "nnnn" / INFO CODE: "yyyy"

TYPE: Error

REASON: An attempt was made to allocate an MMG work data set for a graphic analysis inquiry. The system was unable to successfully allocate the data set.

ACTION: Run terminates with a user 998abend. Check the spelling of the data set name. If the name is spelled correctly, and the data set exists and is cataloged, there can be a more serious system problem. If so, contact your CA MICS system administrator.

REFERENCES: Section 2.7.6.3, Workload Groupings

```
+-----+
| P E R 0 2 0 2 6 |
+-----+
```

TEXT: DYNAMIC ALLOCATION FAILURE ATTEMPTING TO FREE
DDNAME: "ddname", / ERROR CODE: "nnnn" / INFO
CODE: "yyyy". PROCESSING WILL CONTINUE.

TYPE: Warning

REASON: An attempt was made to free the MMG work data set that was previously allocated dynamically. Since the original allocation was successful, this error does not affect the processing of the current inquiry, and it is allowed to proceed.

ACTION: You can check that the MMG workload data set is still accessible. No other action is required.

REFERENCES: Section 2.7.6.3, Workload Groupings

```
+-----+
| P E R 0 2 0 2 7 |
+-----+
```

TEXT: THE LINES PER PAGE WAS SET TO LESS THAN 40, THIS
REPORT REQUIRES A MINIMUM PAGE SIZE OF 40, PAGE
SIZE SET TO 40.

TYPE: Information

REASON: Your MICF options specify a page size of less than 40 lines per page, but this report requires at least 40 lines per page for acceptable output. The page size has been set to 40 lines for this report.

ACTION: Change your MICF options to use a page size of at least 40.

REFERENCES: None

```
+-----+  
| P E R 0 3 0 0 0 |  
+-----+
```

TEXT: GENERIC INFORMATIONAL MESSAGE

TYPE: Information

REASON: This message is used for many reasons.

ACTION: No action required.

REFERENCES: None

```
+-----+  
| P E R 0 3 0 1 0 |  
+-----+
```

TEXT: GENERIC MESSAGE USED TO GENERATE TWO BLANK LINES

TYPE: Information

REASON: Generic message used to generate line spacing on MICSLOG entries.

ACTION: No action required.

REFERENCES: None

```
+-----+  
| P E R 0 3 0 1 2 |  
+-----+
```

TEXT: GENERIC MESSAGE USED TO GENERATE ONE BLANK LINE

TYPE: Information

REASON: Generic message used to generate line spacing on

MICSLOG entries.

ACTION: No action required.

REFERENCES: None

```
+-----+  
| P E R 0 3 0 5 0 |  
+-----+
```

TEXT: "PHASE": "_PGMTTL" BEGINNING EXECUTION...

TYPE: Information

REASON: Processing phase _PGMTTL is starting.

ACTION: None.

REFERENCES: None

```
+-----+  
| P E R 0 3 0 6 0 |  
+-----+
```

TEXT: NEUGENTS MODEL "_FUNCT" PHASE BEGINNING...

TYPE: Information

REASON: Neugents technology internal phase _FUNCT starting.

ACTION: None.

REFERENCES: None

```
+-----+  
| P E R 0 3 0 7 0 |  
+-----+
```

TEXT: "PHASE": "_PGMTTL" ENDING EXECUTION...

TYPE: Information

REASON: Processing phase _PGMTTL ending.

ACTION: None.

REFERENCES: None

```
+-----+  
| P E R 0 3 0 7 5 |  
+-----+
```

TEXT: "PHASE": "_PGMTTL" ABNORMALLY TERMINATING

TYPE: Error

REASON: Processing _PGMTTL abnormally terminated.

ACTION: Neugents technology processing job ends.

REFERENCES: None

```
+-----+  
| P E R 0 3 0 8 0 |  
+-----+
```

TEXT: "V1" FILE CONTAINS "N" RECORDS SELECTED FOR
PROCESSING.

TYPE: Information

REASON: N records have been selected for processing from
file V1.

ACTION: None.

REFERENCES: None

```
+-----+  
| P E R 0 3 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 to data value errors or unmatched
observations among files. Modeling results can
be affected and should be examined because of
this message.

ACTION: None.

REFERENCES: None.

```
+-----+  
| P E R 0 3 0 9 0 |  
+-----+
```

TEXT: NO FORECAST FILE ENTRY FOUND FOR "V1" ENTRY:,

TYPE: Warning.

REASON: An observation in the primary PER file was not matched in the alternate input file.

ACTION: The original observation is discarded from further processing. See message PER03095 for additional details.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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.

```
+-----+  
| P E R 0 3 1 0 0 |  
+-----+
```

TEXT: REPORT: "V1" BEING CURRENTLY GENERATED ON DDNAME //"V1"

TYPE: Informational

REASON: Generation of report V1 has started.

ACTION: The report is written on ddname V1. This message is only issued if the MICF environment is not being used. Message PER03105 is normally issued instead.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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 to the SASLIST DD.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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.

```
+-----+  
| P E R 0 3 1 2 0 |  
+-----+
```

TEXT: GRAPH: "V1" BEING CURRENTLY GENERATED ON DDNAME
// "V1"

TYPE: Informational

REASON: Generation of graph V1 has started.

ACTION: The graph is written on ddname V1. This
message is only issued if the MICF environment
is not being used. Message PER03125 is normally

issued instead.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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 to the SASLIST DD.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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.

```
+-----+  
| P E R 0 3 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 discarded from further processing. Report BADTRAIN contains a list of observations that were dropped.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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 is written on ddname MODLSTAT. This message is only issued if the MICF environment is not being used. Message PER03165 is normally issued instead.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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 to the SASLIST DD.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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.

```
+-----+  
| P E R 0 3 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 technology.

ACTION: Neugents technology processing is terminated.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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.

```
+-----+  
| P E R 0 3 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.

```
+-----+  
| P E R 0 3 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 PER03310 or PER03590.

ACTION: Neugents technology processing terminates.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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 Technology 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.

```
+-----+
| P E R 0 3 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 Technology 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.

```
+-----+
| P E R 0 3 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 executes 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 is written on ddname MODLSTAT. This message is only issued if the MICF environment is not being used. Message PER03435 is normally issued instead.

REFERENCES: None.

```
+-----+  
| P E R 0 3 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 executes 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 to the SASLIST DD.

REFERENCES: None.

```
+-----+  
| P E R 0 3 5 1 0 |  
+-----+
```

TEXT: OPEN FAILED FOR "STD" FILE. SYSTEM ERROR
MESSAGE IS "SYSMSG" EXECUTION TERMINATED.

TYPE: Error

REASON: An open was issued for files with ddname STD. Such files are used by Neugents technology processing for passing data between Neugents technology and the CA MICS Performance Manager

Option 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 causes processing to be terminated.

ACTION: Review the specific system message for action.

REFERENCES: None.

```
+-----+  
| P E R 0 3 5 9 0 |  
+-----+
```

TEXT: ERROR TEXT FROM NEUGENTS PROCESSING
"ERRMSG" EXECUTION TERMINATED.

TYPE: Error

REASON: Any error text returned by USS in the STDERR data set is displayed herein.

ACTION: Review specific error message for action.

REFERENCES: None.

```
+-----+  
| P E R 0 3 6 0 0 |  
+-----+
```

TEXT: DATA FOUND IN THE BATMUG FILE FOR SYSID: LP01
-NO CORRESPONDING DATA WAS FOUND IN THE HARCPU
FILE. REPORTING WILL EXCLUDE SYSID: LP01

TYPE: Warning

REASON: Product usage data was found in the BATMUG file for the given SYSID, but no corresponding SYSID found in the HARCPU file. The problem may be caused because the systems data may not be processed into the RMF unit selected for reporting.

ACTION: Verify that all SYSIDs represented in the BATMUG file are processing into the RMF unit selected for reporting.

REFERENCES: None.

```
+-----+  
| P E R 0 3 6 0 5 |  
+-----+
```

TEXT: SYSID: LP01 OBSERVATIONS FOUND IN THE HARCPU
FILE WITH NO CORRESPONDING DATA IN THE BATMUG
FILE. - OBSERVATION WILL BE DROPPED FROM
ANALYSIS.

TYPE: Informational

REASON: Data was found in the HARCPU file that did not
correspond with any BATMUG data based on the
merge by of SYSID YEAR MONTH DAY HOUR. This
indicates that data from a given day or hour is
missing from the HARCPU file.

ACTION: Verify that the HARCPU file cycle retention is
equal to the number of cycles kept for the
BATMUG file.

REFERENCES: None.

```
+-----+  
| P E R 0 3 6 1 0 |  
+-----+
```

TEXT: SYSID: LP01 OBSERVATIONS FOUND IN THE BATMUG
FILE WITH NO CORRESPONDING DATA IN THE HARCPU
FILE. - OBSERVATION WILL BE DROPPED FROM
ANALYSIS.

TYPE: Informational

REASON: Data was found in the BATMUG file that did not
correspond with any HARCPU data based on the
MERGE BY of SYSID YEAR MONTH DAY HOUR. This
indicates that data from a given day or hour is
missing from the BATMUG file.

ACTION: Verify that the BATMUG file cycle retention is
equal to the number of cycles kept for the
HARCPU file.

REFERENCES: None.

```
+-----+  
| P E R 0 3 6 2 0 |  
+-----+
```

TEXT: SYSID: XE09 FROM THE HARCPU FILE HAS MISSING
 VALUES FOR CPUAVLAC.
 - DATA WILL BE INCLUDED IN REPORTING, BUT WILL
 NOT CONTRIBUTE TO THE TOTAL SYSTEM USAGE.

TYPE: Informational

REASON: Data was found in the HARCPU file that has
 missing values for the Long-Term Average of CPU
 Service in MSUs (CPUAVLAC). There can be
 several reasons for this element to have
 missing values. See ACTION below to begin.

ACTION: Review the data dictionary for a detailed
 explanation of CPUAVLAC and review its settings
 in RMFGENIN.

REFERENCES: None.

```
+-----+  
| P E R 0 3 6 3 0 |  
+-----+
```

TEXT: NO DATA FOUND IN BATMUG THAT MATCHED YOUR
 VENDOR SETTING OF:" vendor name"
 QUERY TERMINATING ...

TYPE: Error

REASON: No data selected from the BATMUG file had a
 Vendor/Owner (MUOWNER) equal to the Vendor
 setting on the query composition panel.

ACTION: Remove the Vendor setting on the composition
 panel. Turn on the frequency reports and turn
 off the other reports. Execute the query and
 review the "Frequency Report of MUOWNER:
 Vendor/Owner" for a list of vendors represented
 in your data. Note the Vendor you wish to
 report on and reset the Vendor setting on the
 composition panel with the correct vendor name.
 Activate the reports you want to see and submit
 the query.

REFERENCES: None.

```
+-----+  
| P E R 0 3 6 3 5 |
```


+-----+

TEXT: SYSID: LPAR OBSERVATIONS FOUND IN THE BATMUG
 FILE WITH MORE THAN 16 CONFIGURED CPU.
 THE 16TH CPU CAPABILITY ADJ. FACTOR, CAPMAF16,
 WILL BE USED IN THE SU CONVERSION FACTOR
 CALCULATION.

TYPE: Informational

REASON: Although it is possible to have more than 16
 configured CPs, the SMF-89 record only provides
 CPU adjustment factors for up to 16 processors.

ACTION: In case more than 16 configured processors are
 encountered in the data, the CPU adjustment
 factor for 16 processors will be used in the SU
 conversion factor calculation.

REFERENCES: None.

+-----+

| P E R 0 3 8 0 0 |

+-----+

TEXT: FILE/TIMESPAN VALIDATION ERROR: The "FILE"
 file must be active at the "TIMESPAN" timespan.
 Review the GENLIB settings for the "FILE" file.

TYPE: Error

REASON: The CA MICS file used in this query was not
 active at the given timespan.

ACTION: Review the GENLIB settings for the indicated
 file.

REFERENCES: None.

+-----+

| P E R 0 3 9 0 0 |

+-----+

TEXT: CA MICS Performance Manager Threshold
 Reporting And Alerting Initiated.

MSU Consumption Rate and Soft Capping Analysis
Threshold Exceptions.

Soft Capping percentage has exceeded 07% for the following LPARs in one or more reporting intervals:

SYSplex	SYSID	LPAR_NAME	EXCEPTION_DATA_TIME
PLEX0A	MJ01	MJ01	15JUNyyy:16:14:00
PLEX0B	MJ02	MJ02	15JUNyyy:16:14:00

TYPE: Informational

REASON: The Soft Capping percentage threshold set in the extended options has been exceeded for one or more LPARs in one or more reporting intervals.

ACTION: Look for the eye catcher "<- EXCEEDED 05%" in the MSU Consumption Rate and Soft Capping Report to see all the LPARs and time intervals that exceeded the specified threshold.

If Threshold Alerting is active, see the MICSLOG message BAS63099 to verify that the alert was generated successfully.

REFERENCES: None.

```
+-----+
| P E R 0 3 9 2 0 |
+-----+
```

TEXT: Defined Group Capacity Reports Not Generated.
No LPARs found with a Defined Group Capacity.

TYPE: Informational

REASON: The Defined Group Capacity Reports can not be generated because there were no LPARs in the input data source with a defined group capacity.

ACTION: Confirm that the SYSIDs used as input have a group capacity defined.

REFERENCES: None.