LC28-1556-2
File No. S370-34

MVS/Extended Architecture Resource Measurement Facility (RMF) Monitor I and II Reference and User's Guide

Program Product

Program Number 5665-274

IBM

LC28-1556-2
File No. S370-34

MVS/Extended Architecture Resource Measurement Facility (RMF) Monitor I and II Reference and User's Guide

Program Product

Program Number 5665-274

IBM

Third Edition (September, 1989)

This is a major revision of, and obsoletes LC28-1556-1 and Technical Newsletter LN28-1189. See the Summary of Amendments following the Contents for a summary of the changes made to this manual. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

This edition applies to Version 3, Release 5.1 of the program product RMF (Resource Measurement Facility) and to all subsequent releases until otherwise indicated in new editions or Technical Newsletters. Changes are made periodically to the information herein; before using this publication in connection with the operation of IBM systems, consult the latest *IBM System/370 Bibliography,* GC20-0001, for editions that are applicable and current.

Note to all users:

Expanded storage is also known as extended storage. The terms are used interchangeably throughout the RMF Library. This book uses the term extended storage.

About This Book

The Resource Measurement Facility (RMF) is a measurement collection tool designed to measure selected areas of system activity and present the data collected in the form of SMF (System Management Facility) records, formatted printed reports, or formatted display reports. An installation can use the information in the RMF output to evaluate system performance and identify reasons for performance problems.

This book provides a description of what RMF is, what it can do, and how to use the RMF Monitor I and Monitor II sessions. For information on the Monitor III session, see *MVS/XA Resource Measurement Facility (RMF) Monitor III Reference and User's Guide*. For information on the System Availability Management function of RMF, see *MVS/XA Resource Measurement Facility (RMF) System Availability Management User's Guide*.

The book contains three types of information:

- A basic, conceptual description of what RMF is and how it can help you to measure system performance.
- Procedural information that describes step-by-step approaches to making RMF work for you.
- Reference information that enables you to locate a specific piece of required information.

Trademarks

The following are trademarks of International Business Corporation.

- MVS/XA™
- MVS/SP™
- Processor Resource/Systems Manager™
- PR/SM™

Who This Book Is For

This book is intended for use by the system programmer responsible for installing RMF and modifying its functions, by the system personnel responsible for measuring and improving system performance, and by the operator and terminal user. Because RMF is a tool for measuring system performance, this book assumes that the reader has extensive knowledge of the MVS system.

How This Book Is Organized

The contents and suggestions for use of each chapter are:

Chapter 1. Introduction: presents basic, conceptual information about RMF. Everyone involved with installing, operating, modifying, and using RMF should read this chapter.

Chapter 2. RMF Control: presents the system commands, session commands, display commands, and JCL needed to run RMF Monitor I and II session. Separate sections describe RMF initialization, RMF session control, displaying RMF status, RMF termination, and RMF options for Monitor I and II.

Chapter 3. Monitor I Session: describes the options used to control Monitor I session processing and the session reports produced. For each report, the chapter includes suggestions for using the information in the report, a description of each field, and at least one sample report. The reports are described in alphabetical order. Note that each field in the report also appears in the index to help you locate the description for a specific field as rapidly as possible.

Chapter 4. Monitor II Session: describes the options used to control Monitor II session processing and the session reports produced. For each report, the chapter includes suggestions for using the information in the report, a description of each field, and at least one sample report. The reports are described in alphabetical order. Note that each field in the report also appears in the index to help you locate the description for a specific field as rapidly as possible.

Chapter 5. Post Processor: describes SMF data sets and control statements used in controlling the post processor. This chapter also describes the duration, summary, exception, and plot reports the post processor generates. The chapter includes suggestions for using the information in each type of report, a description of each field, and at least one sample report.

Chapter 6. Procedures For Adding User Reports: describes how to add a user report to Monitor I session processing, how to add a field to the fields traced by RMF, how to add a user report to Monitor II session processing and how to add a routine to limit the use of RMF under Monitor II TSO.

Chapter 7. SMF Records: briefly describes each SMF record built by RMF. The records are fully described in *MVS/Extended Architecture System Management Facilities (SMF)*.

Chapter 8. Fixed Storage Requirements: consists of a table that you can use to determine, according to the options you specify, the fixed storage requirements for the Monitor I session.

Chapter 9. Resource Measurement Facility User Completion Codes: lists, in numerical order, each user completion code RMF issues for Monitor I and II. Problem determination information is also provided. To help the operator to respond to the messages RMF issues, you might want to place this chapter with the system completion code information kept at the operator console.

Related Information

To use this book effectively, you need to know about the information in the following book:

Figure 1-1. RMF General Information

Short Title Used in This Book	Title	Order Number
RMF General Information Manual	*Resource Measurement Facility (RMF) Version 3 General Information Manual*	GC28-1115

While you are using this book, you may need the information in the following books:

Figure 1-2. RMF Publications

Short Title Used in This Book	Title	Order Number
RMF Monitor III User's Guide	*MVS/Extended Architecture Resource Measurement Facility (RMF) Monitor III Reference and User's Guide*	LC28-1557
RMF Messages and Codes	*MVS/Extended Architecture Resource Measurement Facility (RMF) Messages and Codes*	GC28-1382
RMF Operating Procedures for Monitor I and II	*MVS/Extended Architecture Resource Measurement Facility (RMF) Operating Procedures: Reference Summary for Monitor I and II*	SX22-0009
RMF Operating Procedures for Monitor III	*MVS/Extended Architecture Resource Measurement Facility (RMF) Operating Procedures: Reference Summary for Monitor III,*	SX22-0010
RMF Programming Logic Manual Volumes I and II	*MVS/Extended Architecture Resource Measurement Facility (RMF) Programming Logic Manual, Volumes I and II,*	SX22-0010
RMF SAM User's Guide	*MVS/Extended Architecture Resource Measurement Facility (RMF) System Availability Management User's Guide*	SC28-1558
RMF SAM Diagnosis Guide	*MVS/Extended Architecture Resource Measurement Facility (RMF) System Availability Management Diagnosis Guide*	LC28-1559

Where necessary, this book references information in other books, using shortened versions of the book title. The following table shows the shortened titles, complete titles, and order numbers of the books that you might need while you are using this book.

Figure 1-3. Related Publications

Short Title Used in This Book	Title	Order Number
SMP System Programmers Guide	*OS/VS System Modification Program (SMP) System Programmer's Guide*	GC28-0673
MVS/XA Initialization and Tuning	*MVS/Extended Architecture System Programming Library: Initialization and Tuning*	GC28-1149
MVS/XA System Management Facilities	*MVS/Extended Architecture System Programming Library: System Management Facilities (SMF)*	GC28-1153
MVS/XA System Messages Volume I	*MVS/Extended Architecture Message Library: System Messages, Volume I*	GC28-1376
MVS/XA System Messages Volume II	*MVS/Extended Architecture Message Library: System Messages, Volume II*	GC28-1377
MVS/XA System Codes	*MVS/Extended Architecture Message Library: System Codes*	GC38-1157
MVS/XA Routing and Descriptor Codes	*MVS/Extended Architecture Message Library: Routing and Descriptor Codes*	GC28-1194.
MVS/XA Utilities	*MVS/Extended Architecture Utilities*	GC26-4018
MVS/XA Service Aids	*MVS/Extended Architecture System Programming Library: Service Aids*	GC28-1159
MVS/XA System Modifications	*MVS/Extended Architecture System Programming Library: System Modifications*	GC28-1152
MVS/XA TSO Guide	*MVS/Extended Architecture TSO Guide to Writing a Terminal Monitor Program or a Command Processor*	GC28-1295
OS/VS2 TSO Guide	*OS/VS2 System Programming Library: TSO*	GC28-0629

Summary of Amendments

Summary of Amendments
for LC28-1556-2
for RMF Version 3 Release 5.1

This major revision, which supports Version 3 Release 5.1 of the program product Resource Measurement Facility (RMF), describes changes in RMF for MVS/System Product Version 2 Release 2.3.

The changes include:

- Addition of storage group fields to the Monitor I I/O device activity report for direct access devices
- Addition of storage group fields to the Monitor II I/O device activity row (DEVV) and table (DEV) reports
- Changes to SMF record type 74 and type 79 subtype 9 to include storage group measurements.

Minor technical and editorial changes have been made throughout the book.

Summary of Amendments
for LC28-1556-1
as updated on September 1, 1988
by Technical Newsletter LN28-1223

This newsletter incorporates changes in support of Version 3 Release 5 of the program product Resource Measurement Facility (RMF) for MVS/System Product™ Version 2 Release 2. Editorial, maintenance, and technical changes are also included.

Trace support for variables RCEESREA, RCEVIOME, RCEVIORE, and RCEVIOMG has been added to the Monitor I Trace report.

The System Measurement Facility (SMF) 71 records have been extended to include virtual input output (VIO) page rates.

Summary of Amendments
for LC28-1556-1
as updated on June 17, 1988
by Technical Newsletter LN28-1245

This newsletter incorporates changes in support of Version 3 Release 5 of the program product Resource Measurement Facility (RMF) for MVS/System Product Version 2 Release 2. Editorial, maintenance, and technical changes are also included.

These updates support the Processor Resource/Systems Manager™ (PR/SM™) enhancement, which allows you to configure a 3090 Processor Complex Enhanced model into logical partitions. Each partition can run a different operating system. Changes include a new partition data report section of the CPU activity report, which provides information on all configured partitions.

Contents

Figures

Chapter 1. Introduction

The MVS/XA Resource Measurement Facility (RMF) is a powerful and flexible tool for measuring the performance of your system and helping you to pinpoint the sources of performance problems. Effective system tuning requires a measurement tool that is easy to use, has little impact itself on system performance, and is capable of reporting required data at various levels of detail. RMF is such a measurement tool; using it enables you to take a structured, disciplined approach to managing the performance of your system.

Experience has taught us that a haphazard approach to system tuning seldom results in performance improvements. As a result, a methodology — a series of steps — has emerged; this effective approach to system tuning has much in common with true experimentation. It includes the following processes:

1. Use RMF to measure the current performance of your system in concrete terms, such as service provided, transaction rate, TSO response time, or hardware and software availability.
2. Define the objectives for your system in the same concrete, measurable context.
3. Use the measurements obtained during the first step to determine the general areas, such as I/O, processor, or real storage, where constraints are occurring, or when system or application outages are occurring. Based on these high-level measurements, use RMF to take more detailed measurements to determine the reasons for the constraints, enabling you to direct your tuning efforts toward specific problem areas.
4. Once these areas are identified, determine what steps to take to improve performance. Take these steps one at a time to isolate the effect of each improvement. To determine the effectiveness of the step you have taken, evaluate the result and compare the measurement to the earlier measurements and the concrete objectives you have set. Repeat the process until your objectives are met.
5. Continue to measure and evaluate your system's performance to enable you to spot performance problems before they become critical.

The common factor in all of these steps is constant measurement of the system's performance, and RMF provides you with both the ability to constantly measure the high-level performance indicators and the ability to take detailed measurements of many problem areas indicated by the high-level measurements. The data obtained can help you to:

- Identify system components that show patterns of exceptional utilization.
- Identify periods of system activity during which the utilization of particular resources is exceptional.
- Relate the service different classes of users are receiving to the specification provided in the IPS.
- Identify bottlenecks of resource contention.
- Locate excessive users of particular system services.
- Identify periods of system and application failure (outages).

One way to use RMF is to constantly measure key areas of system activity (such as processor activity, paging activity, virtual storage activity, workload activity, channel path activity, I/O queuing activity, and direct access device activity) through a Monitor I session and record the data gathered in SMF records. These records, when archived, provide machine readable performance data over an extended period of time. You can run the RMF post processor at regular intervals to report the performance data. You can limit the volume of printed reports by requesting duration or summary reports, which summarize system activity, or by using exception reporting, which produces reports only when an exceptional condition occurs.

A second way to use RMF is at a display station to constantly measure the workflow, or speed, with which jobs move through the system and note any exceptions to conditions you specify. RMF gathers data about system performance and reports the data while the system is running. The reports present the data in a format that simplifies the process of identifying a potential performance problem.

Note: There are two kinds of report sessions that can display sampled data: Monitor II and Monitor III sessions. Monitor III sessions are described in *Resource Measurement Facility (RMF) Monitor III Reference and User's Guide*, LC28-1557.

A third way to use RMF is through means of the System Availability Management (SAM) function. RMF SAM provides the means to record, track, and report software and hardware availability. Working with Information Management, SAM can identify problems that involve systems or application failures (outages) and generate user-defined reports. You can tailor reports to present information about various hardware and software problems through ISPF panels. This information can be used by the SAM report writer to produce reports.

Note: For information about RMF SAM, see *RMF System Availability Management User's Guide*, LC28-1558.

To allow you to choose how you can best measure the performance of your system, there are different ways to execute RMF. Each execution of RMF is called a session, and, as with the resources that can be measured and the types of output available, you choose the type of session that best meets your needs. During an RMF session, the basic unit of measurement and reporting is called an interval. At the beginning of an interval, RMF initializes to measure data. At the end of an interval, RMF summarizes the data it has measured, formats the data, and reports it in the form you have selected.

RMF measures data in two ways: by exact count and by sampling. RMF obtains data that is measured by exact count by checking the appropriate system indicators at the beginning and the end of the interval and computing the difference.

RMF obtains data that is sampled by checking the appropriate system indicators at each cycle within an interval. A cycle is a unit of time within an interval, usually relatively small in comparison to the length of the interval. At each cycle, RMF samples the data it requires. At the end of the interval, the mass of data collected at each cycle is reduced, if necessary, for reporting—for example, to present minimum, maximum, and average values for a field or to present data as a percentage.

Note: When the sampling technique is used, a random independent distribution of events is assumed.

You choose the length of the interval for each type of RMF session and the length of the cycle for the type of RMF session that can sample data. You can also synchronize the interval to any time within the hour to allow easy comparison of data from different days or different systems.

The following sections describe the resources RMF can measure, the types of output it can generate, and the types of sessions that are available. This chapter also includes information on system requirements for running RMF.

Measurement Activities

RMF can measure resource utilization in the following areas of system activity:

- Processor (CPU) activity, enabling you to determine how much wait time each processor experiences.
- Address space activity, in a variety of forms. The processor activity report includes an address space analysis section that gives an overview of how address spaces are being used and their status. Additional reports, available either at a display station or in printed or SMF record form, provide more detailed information. These reports cover address space state data, address space resource data, and address space SRM data. You can obtain these reports for all address spaces, groups of address spaces, or individual jobs. Thus, you can use RMF to obtain address space activity analysis at whatever level of detail you require.
- Channel path activity provides information on channel path use for all channel paths in the system. This information, when used in conjunction with the I/O device activity and I/O queuing activity information (or with trace information for logical path utilization), can give you the information you need to evaluate your channel path configuration.
- I/O device activity can be provided for individual devices or for a device class. I/O device measurements identify overall response time for the device, the average time spent in each stage of an I/O request (IOS queue time, pending time, connect time, and disconnect time), and percentages that indicate device use. RMF can measure I/O device activity for several device classes: direct access devices, direct access devices within a storage group, magnetic tape devices, unit record devices, communication equipment, graphics devices, and character reader devices. You choose the classes or specific devices to be measured, depending on your particular configuration and needs.

  Individual device measurements can be provided for all devices in one or more device classes or limited to one or more specific devices in the system. In a report, the devices are grouped by logical control unit (LCU) and/or by storage group (SG).
- I/O queuing activity provides a different set of I/O measurements for the devices in a logical control unit group. The definition of a logical control unit (LCU) is model-dependent.

  On a 308x, or 908x processor, an LCU is the set of devices attached to the same physical control unit (or group of control units that have one or more devices in common). Each device belongs to only one LCU and all requests to the various devices within an LCU share the same LCU queue for I/O request initiation and redrive within the channel subsystem.

  On a 4381 processor, an LCU is the set of devices attached to an identical set of physical control units. A physical control unit can belong to more than one LCU if the devices are configured asymmetrically. Each device belongs to only one LCU. All I/O requests and redrives are queued on a physical control unit basis.

  On a 3090 processor, an LCU is the set of devices attached to the same physical control unit (or group of control units that have one or more devices in common). As with the 308x processor, each device belongs to only one LCU, but the I/O processor (IOP), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the various LCUs of a 3090 processor.

Logical control unit queue length (for the 308x/908x and the 3090 processor only) and I/O request delay reasons are provided, along with online channel path/control unit configuration information to help with configuring the I/O subsystem.

RMF obtains the I/O configuration data from the I/O configuration data set (IOCDS) for a 308x/908x or a 3090 processor, or, for the 4381 processor, from the I/O configuration data (IOCD) in the hardware system area (HSA).

- Paging activity, which can be reported on two levels. RMF provides a system overview report that you can obtain at a display station for immediate inspection. RMF also provides a paging activity report that gives detailed information on paging and swapping activity. The report can be printed, and the data can be stored in an SMF record. These reports can alert you to a performance problem caused by excessive paging or swapping and give you the information you need to pinpoint the cause of the problem and take steps to correct it.
- Workload activity, which can be reported at three levels: performance group period, performance group, and system summary. By assigning unique performance group numbers and domain numbers to a particular user or set of users (such as subsystems, monitor programs, groups of users, and batch initiator classes), you can use the workload activity reports to show the system services provided for the particular user or set of users.
- Page/swap data set activity, which includes separate sections for page data sets and swap data sets. The information in these reports can help you to verify that the space allocation for each data set is right for your particular installation.
- Virtual storage activity, which always provides a summary report on the use of common storage and can provide a summary report on the use of private area storage for a specific long-running job. These reports help your installation to determine its current use of virtual storage and provide information you can use to predict the effect of future application growth. Detail reports for both common storage and the private area provide additional information.
- ASM/RSM/SRM trace activity, which enables you to trace the contents of certain fields in the RCE control block, ASMVT control block, SRM data area, and SRM domain tables and the SMF data area (SMCA).

  Many of the fields RMF can trace are used as input to RMF reports, where the data is summarized for the entire interval. Tracing such a field gives you a more detailed picture of how the contents of the field change over the interval. Other fields that RMF can trace contain data that is useful only in specific situations.

  One important use of RMF tracing is to monitor the SRM multiprogramming level (MPL) adjustment. Monitoring the MPL adjustment can enable you to determine if the domain setting in the installation performance specification (IPS) is correct.

  There is a user exit provided in the tracing routine that enables you to add a field or fields to those offered by RMF.
- Enqueue activity, which measures the contention for or ownership of serially-reusable resources that are controlled by enqueuing. Enqueue activity can be reported as an interval report or as a display report, available for immediate inspection. Both the interval report and the display report have a

summary level and a detail level; the detail report includes the names of jobs that are contributing to the contention for a particular resource.

RMF also provides a display report that measures contention caused by reserved shared direct access devices (shared DASDs).

The information on enqueue and reserve activity can be of great help in spotting points when excessive contention is occurring. You can then take steps, such as balancing the workload, to minimize performance problems caused by jobs waiting on a particular resource.

- Real storage/processor/SRM activity, which gives a one-line summary of the use of real storage, the processor, and SRM facilities. The report is available at a display station for immediate inspection. You can use the display report to alert you to the beginning of a performance problem so that you can initiate more detailed reporting to obtain information about the problem as it develops.
- Transaction activity, which gives an overview of transaction activity by performance group period, is available at a display station. The report enables you to evaluate the effectiveness of the performance objectives in the IPS and the workload classification in the installation control specification.
- Domain activity, which contains the same information as the system response to the DISPLAY command with the DMN operand. The report is available at a display station. The report enables you to evaluate the effectiveness of the current domain definitions.

Your installation can measure other areas of system activity by replacing the user exit routines RMF provides with your own routines that measure and report on an activity that your installation requires.

Types of Output

System activity measurements are available in three forms:

- printed reports
- SMF records
- display reports

RMF can report all activity measurements except workflow and delay as SMF records and/or printed interval reports. SMF records are written at the end of each measurement interval. Interval reports can be generated at the end of each interval, at the end of a session, at the end of RMF processing, or at a later time by the post processor.

The post processor can generate additional reports. It creates these reports from data collected during a Monitor I session over a reporting period that normally spans multiple intervals. Post processing enables your installation to generate reports during off shift hours and is especially useful when you are using RMF as a high-level measurement tool. The additional post processor reports are:

- Duration reports, which can summarize processor activity, channel path activity, I/O device activity, I/O queuing activity, paging activity, page/swap data set activity, workload activity, and virtual storage activity over a reporting period that you select.

- Summary reports, which can contain one line of data summarizing system activity during each interval included in the reporting period and a single line of data containing total summary figures for all intervals in the reporting period.
- Exception reports, which provide an overview of system activity, like a duration report or a summary report. However, unlike a duration or summary report, an exception report does not average or summarize data over a period of time; an exception report occurs only when a key measurement that you select exceeds a threshold value that you supply for that measurement.

  When an exception occurs – when a key measurement exceeds its threshold – the post processor generates an exception report that includes a line of data for each exception detected.
- Plot reports, which present a graphic view of many areas of system activity, such as the time the processor is busy, the number of batch users, or the activity rate for a channel path or device.

Display reports, which provide immediate access to measurement data, are generated during a Monitor II display session. You can also obtain printed copies of either single reports or all reports during a display session. RMF can provide display reports for most system activities (paging, address space, channel path, device, I/O queuing, domain, enqueue, page/swap, and real storage/processor/SRM).

You can choose any or all of the forms of reported data to give your installation the most useful information. The reported data can then be used to tune the system for better performance and to plan for resource changes.

Types of Sessions

To give flexibility in the types of system activity you can measure and the type of output you can choose, RMF executes during several types of sessions. Although the term "session" is normally used to refer to an interactive use of computing facilities, an RMF session is, simply, an execution of RMF.

Two types of RMF sessions are described in this manual:

- Monitor I
- Monitor II

For information about Monitor III sessions and reports, see *MVS/XA Monitor III Reference and User's Guide.*

Monitor I and Monitor II sessions collect and report data about specific system activities.

Monitor I Session

A Monitor I session is a continuous background session of any duration that collects information about processor activity, channel path activity, I/O device activity, I/O queuing activity, workload activity, virtual storage activity, paging activity, enqueue contention activity, page and swap data set activity, and ASM/RSM/SRM trace activity. User exits allow you to measure and report on other areas of system activity that are relevant to your installation. Data reported from the activities measured can be in the form of SMF records and printed reports. Either or both of these types of output can be selected.

The activities to be measured, the type of output required, the length of the interval, the length of the sampling cycle, and the length of the session itself are specified on options that you supply to RMF at the beginning of the session. An operator command starts a Monitor I session. Only one Monitor I session can be active at any particular time, but you can modify the options during the session. Thus, although the Monitor I session has potentially long duration, the activities it measures and the type of output can be changed at any time during the session to meet changes in your information needs. An operator or session command stops the Monitor I session.

Monitor II Session

A Monitor II session is basically a snapshot session; that is, unlike the ongoing Monitor I session that can measure various areas of system activity over a long period of time, a Monitor II session generates a requested report from a single data sample. A Monitor II session running independently of the Monitor I session can collect information about address space activity, reserve activity, channel path activity, enqueue activity, real storage/processor/SRM activity, and domain activity. A Monitor II session running at the same time as a Monitor I session can produce reports that overlap some areas of Monitor I session measurements. These reports and the corresponding Monitor I session measurements are:

Monitor II Session Reports	Monitor I Session Measurements
transaction activity	workload activity
device activity	device activity
I/O queuing activity	I/O queuing activity
page or swap data set activity	page /swap data set activity

The types of output for the resources measured during a Monitor II session can be in the form of reports sent to a display station for immediate inspection or in the form of SMF records and/or printed reports. To allow for this variety of output forms, Monitor II sessions can be either display sessions or background sessions. You can have as many TSO sessions as you require. However, only 32 non-TSO Monitor II sessions can be active at any one time.

To enable you to take a snapshot of other areas of system activity during either a display or a background session, you can add your own data gathering and reporting routines to Monitor II session processing.

A Monitor II Display Session: is either a task that runs as a local 3270 display session or a TSO display session. Both types of sessions measure the same areas of system activity and generate output in the form of display reports for immediate inspection, with the option of obtaining printed copies of individual reports or all reports produced during a session. The differences between the two types of display sessions lie in the area of operating techniques:

> **Local 3270 Display Session:** A local 3270 display session is started by an operator command. Once started, it is controlled by the terminal user through display commands. The session can be stopped by either the operator or the terminal user, but only the terminal user can modify its processing. Because the session is started by an operator command, RMF must be initialized before the session can be started.

TSO Display Session: A TSO display session is started when a terminal user who has logged on to TSO issues the RMFMON command. Once started, the terminal user controls the session through display commands. The session can be stopped or modified only by the terminal user. A TSO display session does not require that RMF be initialized.

A Monitor II Background Session: measures the same areas of system activity as a display session, but it generates output in the form of printed reports and/or SMF records. Either or both output forms can be selected for a particular session. A background session is started by an operator command, and processing, such as the interval between the snapshot reports, is controlled by options determined when the session is started. These options can be modified during the session. The session is stopped by an operator command or the expiration of a user-specified length of time for the session. Because the session is started by an operator command, RMF must be initialized before the session can be started. Because the types of system activity measured and the types of output generated differ for each session, the type of session required depends on what you need to know about your system's performance at any particular time and what form of output you require.

In addition to the wide choice of resource measurement and resource reporting offered by the various types of sessions, the post processor can generate interval, duration, summary, plot reports and exception reports whenever convenient. You can thus use RMF in a way that best meets your installation's particular needs and procedures.

Chapter 2. RMF Control

This chapter describes the procedures needed to execute an RMF Monitor I or Monitor II session.

You use commands and JCL statements to initialize, control, modify, and terminate an RMF session. This section includes the following topics:

- "RMF Commands" describes the three categories of commands that control RMF: system commands, session commands, and display commands.
- "RMF JCL" describes the JCL statements RMF requires.
- "RMF Initialization" tells how to initialize RMF by using the START system command or by submitting a batch job.
- "RMF Session Control" describes how to identify and control a particular RMF session by using session commands.
- "Displaying RMF Status" describes how to determine what sessions are active and what options are in effect for each session.
- "RMF Termination" discusses how to stop RMF processing by using the STOP system command.
- "RMF Options" describes how to tailor an RMF session to suit your particular needs.

Examples of commands are used to illustrate each RMF control function. Some of these commands include Monitor I session options. The full meaning of each option is described later under "Monitor I Session Options" in Chapter 3.

RMF Commands

The commands that control Monitor I and II sessions fall into three categories: system commands, session commands, and display commands.

System Commands

There are three system commands that control RMF processing: START, MODIFY, and STOP.

START: The START system command initializes RMF or both initializes RMF and starts a Monitor I session.

MODIFY: The MODIFY system command passes **session commands** to RMF sessions.

STOP: The STOP system command terminates all RMF sessions (except those running under TSO).

Session Commands

The four session commands that control RMF Monitor I and II session processing are START, MODIFY, STOP and DISPLAY. The operator uses the input field of the **system** MODIFY command to pass any of these commands to RMF.

START: The START session command starts the Monitor I session, a Monitor II background session, or a Monitor II local 3270 session. The operator must issue the START **system** command to initialize RMF before issuing a START **session** command.

MODIFY: The MODIFY session command changes the options currently in effect for an active Monitor I or Monitor II background session. This command causes RMF to end the current sampling interval, change any options specified in the command, and resume processing. Any changed options take effect immediately.

STOP: The STOP session command terminates the Monitor I session, a specific Monitor II background session, or a specific Monitor II local 3270 display session.

DISPLAY: The DISPLAY session command displays status information for active RMF sessions (except those running under TSO). Depending on the options on the command, you can display the session identifiers of all active non-TSO sessions, display the options for a particular session, or display the session identifiers and options for all active non-TSO sessions.

Display Commands

The terminal user controls RMF processing for a session that produces display station output (Monitor II local 3270 and TSO display sessions). The terminal user issues display commands and selects reports. Display commands perform many of the same functions performed by option selection during a Monitor I session, Monitor II background session, or Monitor III session.

Monitor II display commands enable the terminal user to:

- Display the menu -- that is, the list of reports that are available.
- Select an item from the menu and cause the selected report to be displayed once. During a display session on a device that has program function keys (PF keys), you can select an item from the menu by pressing the PF keys that RMF associates with the menu item.
- Select an item from the menu and cause the display of the selected report to be updated automatically at a timed interval.
- Recall the most recently displayed data for a selected report.
- Frame backward and forward through a multi-frame report.
- Direct the current screen image to SYSOUT.
- Direct all reports requested during the session to SYSOUT.
- Specify whether the data presented in the requested reports is to be total data or "delta" data -- data that reflects only the changes since the previous request for a report.
- Stop the display session.
- Change the menu item associated with a PF key .

See Chapter 4 for a detailed discussion of Monitor II display commands.

Syntax Notation

This book uses the "brackets and braces" syntax language to describe the commands, options, and menu items. The set of symbols listed below define the format of each item; you should never enter these symbols in the actual command, option, or menu item.

Underscore: The underscore (_) indicates a default option. If you want to select an underscored alternative, you need not actually specify it when you enter the command, option, or menu item, although, of course, you can.

Braces: Braces { } group related items, such as alternatives, and indicate that you must choose **one** of the items enclosed within the braces, unless you are accepting a default value.

Brackets: Brackets [ ] can also group related items; however, everything within the brackets is optional and can be omitted.

You should use uppercase letters, commas, and parentheses exactly as shown. Lower case letters represent variables for which you should substitute specific information. Stacked items, enclosed in either braces or brackets, indicate alternatives. One or more of the alternatives can be selected.

RMF JCL

Whether you initialize RMF by an operator START system command or as a batch job, certain JCL statements are required.

RMF supplies a cataloged procedure for Monitor I sessions and Monitor II background and local 3270 display sessions that are invoked in response to a START command for RMF or an EXEC statement specifying PROC=RMF or RMF. The contents of the IBM-supplied procedure needed to initialize RMF for these sessions are:

```
//IEFPROC    EXEC  PGM=ERBMFMFC,REGION=4096K,DPRTY=(10,10)
//IEFPARM    DD    DDNAME=IEFRDER
//IEFRDER    DD    DSN=SYS1.PARMLIB,DISP=SHR
//SYSUDUMP   DD    SYSOUT=A
```

Notes:

1. When you are running a large number of sessions, you might find it necessary to override the REGION parameter on the EXEC statement to provide a larger region size.
2. When work with a higher dispatching priority is executing, RMF is not dispatched when its interval time expires. Data collection might be skipped for the interval. To solve this problem, override the DPRTY parameter on the EXEC statement to provide a higher dispatching priority for RMF.

The RMF partitioned data set is normally SYS1.PARMLIB. The contents of the four parmlib members (for Monitor I and II sessions) supplied by IBM are described later in this section under "Contents of the Parmlib Members."

RMF dynamically allocates all message and report data sets to SYSOUT. However, if the system returns a code that indicates that the ddname is in use, RMF proceeds as if the dynamic allocation had been successful. Thus, your installation can allocate the data sets for RMF in the JCL, enabling you, for example, to route output data to a permanent data set rather than to SYSOUT.
RMF data sets can thus be preallocated. The message and report data sets that RMF uses for Monitor I and II sessions and the ddnames for these data sets are:

Ddname	**Use**
MFMESSGE	Allocated during a Monitor I or Monitor II background session, this data set contains messages that do not pertain to a particular session. One MFMESSGE data set is allocated each time RMF is initialized. **Note:** The SYSOUT class for this data set cannot be changed by the SYSOUT (class) parameter in the RMF options; to change the SYSOUT class, you must preallocate the data set.
RMFSCxx	Allocated during a Monitor I or Monitor II background session, this data set contains messages that pertain to a particular session. One data set is allocated for each session, and **xx** is the session identifier. For example, RMFSCZZ is always the ddname for the Monitor I session, and RMFSCDD would be the ddname for a Monitor II background session with a session identifier of DD.

Note: The SYSOUT class for this data set cannot be changed by the SYSOUT (class) parameter in the RMF options; to change the SYSOUT class, you must preallocate the data set.

MFRnnnnn — Allocated during a Monitor I session, this data set contains report output. A ddname is generated and one data set is created for each interval during the session; **nnnnn** is a decimal number from 00001 to 99999.

The first interval is assigned the ddname MFR00001, the second, MFR00002, and the process of generating ddnames is repeated for each successive interval. Thus, if a session has fifteen intervals, ddnames MFR0001 through MFR00015 are used for the interval reports produced.

MFEnnnnn — Allocated during a Monitor I session, this data set is used by ESTAE to re-allocate the report data sets after a recoverable abnormal termination. The meaning of **nnnnn** and the process of assigning ddnames are the same as for MFRnnnnn.

RMFDMxxx — Allocated during a Monitor II local 3270 or TSO display session, this data set contains hardcopy output generated when hardcopy mode is in effect or when a print command is used. One data set is generated for each session; **xxx** is the session identifier for a local 3270 session or TSO for a TSO session.

RMFxxnnn — Allocated during a Monitor II background session, this data set contains report output. One RMFxxnnn data set is allocated for each report; **xx** is the two-character session identifier, and **nnn** is a decimal number from 001 to 999. A ddname is generated for each report requested. The first report is assigned the ddname RMFxx001, the second, RMFxx002; the process of generating ddnames is repeated for each report. Only one ddname is used for each report, regardless of the number of intervals occurring during the session. If session options are modified to stop and then restart a particular report, a new ddname is generated when the report is restarted. For example, if a session with a session identifier of AB is started, and reports are requested for ASD, SRCS, and SPAG, the three ddnames are RMFAB001, RMFAB002, and RMFAB003. If the session options are modified to NOASD, SRCS, and SPAG, then modified again to re-specify ASD, the ASD reports would be sent to a data set with the ddname RMFAB004.

Note: When you omit the DCB characteristics for the message and report data sets described above, the characteristics used are:

```
DCB=(RECFM=VBA,LRECL=137,BLKSIZE=1693)
```

When you change the DCB characteristics, you cannot change the record format; you must specify RECFM=VBA.

RMF Initialization

RMF must be initialized in order to start sessions (except TSO sessions) or display RMF status. You initialize RMF by issuing a START system command, which invokes the RMF cataloged procedure. RMF can also be initialized as a batch job. When you initialize RMF, you can either simply initialize RMF or you can both initialize RMF and simultaneously start a Monitor I session.

The syntax of the START system command is:

```
{START}  RMF[.id],[ddd],[ser],[parm],[keyword=option]
{  S  }
```

id

A one- to eight-character identifier to be assigned to this initialization of RMF. This is an optional field, but it **must** be specified in order to stop or modify RMF processing through subsequent operator commands.

ddd

The device number or device type (such as 3330) used when you are overriding any corresponding unit specifications on the IEFRDER DD statement in the RMF cataloged procedure. This is a positional parameter; the preceding comma can be omitted only when you are omitting this field and all of the following fields.

ser

The volume serial number used when you are overriding the corresponding volume serial specifications on the IEFRDER DD statement in the RMF cataloged procedure. This is a positional parameter; the preceding comma can be omitted only when you are omitting this field and all of the following fields.

parm

The options for the Monitor I session, if a Monitor I session is to be started concurrently with RMF initialization. To start a Monitor I session using options taken from other sources, omit this field. See "Input Merge for the Monitor I Session" and "Monitor I Session Options" in Chapter 3 for more details. When you specify options, each option has the form:

```
option [(value)]
```

When you specify more than one option, each option must be separated from any other options by a comma, and the entire field must be enclosed in parentheses.

If you do not want to start a Monitor I session concurrently with RMF initialization, you must specify the NOZZ option in this field. When you specify NOZZ, it must be the only option specified. See "Monitor I Session Options" in Chapter 3 for a complete explanation of the NOZZ option.

keyword = option

Any keywords required when you are overriding the data set specifications on the IEFRDER DD statement in the RMF cataloged procedure. The keywords specified must be keywords that are valid on a DD statement. If you have replaced the RMF cataloged procedure supplied by IBM with a procedure that contains symbolic parameters, you can specify in this field any symbolic parameters defined in your procedure. Normally, this field is not used.

Examples

1. To use the START command to initialize RMF without starting a Monitor I session, issue the command:

   ```
   START RMF.A,,,NOZZ
   ```

2. To use the START command to initialize RMF and start a Monitor I session, issue the command:

   ```
   START RMF.A
   ```

   All options for the Monitor I session are taken from other sources (see “Input Merge for the Monitor I Session” in Chapter 3), and the session is assigned a session identifier of ZZ.

3. To initialize RMF and start a Monitor I session as a batch job using the RMF cataloged procedure, the following JCL is required to invoke the RMF cataloged procedure:

   ```
   //RMFJOB   JOB   (accounting information)
   //RMF      EXEC  RMF
   ```

   All options for the Monitor I session are taken from other sources (see “Input Merge for the Monitor I Session” in Chapter 3), and the session is assigned a session identifier of ZZ. The operator can override the default options before session initialization is complete.

4. To use the START command to initialize RMF and start a Monitor I session, specifying a system level workload activity report, setting a cycle length of 500 milliseconds, and taking the other options from other sources, issue the command:

   ```
   START RMF.A,,,(WKLD(SYSTEM),CYCLE(500))
   ```

   The Monitor I session is assigned a session identifier of ZZ.

5. To initialize RMF as a batch job and start a Monitor I session, specifying a system level workload activity report, setting a cycle length of 150 milliseconds, and taking the other options from other sources, use the following JCL:

   ```
   //RMFJOB   JOB   (accounting information)
   //RMF      EXEC  RMF,REGION=4096K,
   //               PARM='WKLD(SYSTEM),CYCLE(500)'
   ```

   The Monitor I session is assigned a session identifier of ZZ.

RMF Session Control

Session commands can start and stop all non-TSO RMF sessions, modify the options for an active Monitor I session, Monitor II background session, or Monitor III data gathering session, and display the status of active non-TSO sessions. Only one Monitor I session can be active at any particular time; up to 32 non-TSO Monitor II sessions can be active concurrently.

To enable RMF to identify a particular session, each session must be assigned a unique session identifier. You use this identifier on all session commands for that particular session; it also appears in all RMF messages that pertain to that session.

For the Monitor I session, the session identifier is always ZZ. If you start the Monitor I session when you initialize RMF, the ZZ session identifier is assigned automatically. If you start the Monitor I session with a START session command, you must specify a session identifier of ZZ. See "Starting the Monitor I Session" in Chapter 3.

For a Monitor II session, you assign the identifier when you start the session. The format required for the session identifier depends on the type of session:

- For a local 3270 display session, the session identifier is the three-digit hexadecimal device number, in the form nnn, of the display station used for the session. See "Starting and Ending a Local 3270 Display Session" in Chapter 4.

  You can specify a monitor-id value of (M2), however, if you do not specify a monitor-id, the default value is M2.

- For a background session, the session identifier is any two-character alphameric field (except ZZ) that you choose. See "Starting a Monitor II Background Session" in Chapter 4.

- The session identifier for a Monitor II TSO display session is always TSO; the identifier is assigned automatically when the session is started.

Session commands are passed to RMF through use of the system MODIFY command. The syntax of the MODIFY command is:

```
{MODIFY}  {[RMF.]id, input}
{ F    }  { jobname,       }
```

id
The identifier assigned to RMF when it was initialized.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

input
A session command that starts, stops, or modifies the options of a non-TSO session. This field is also used to issue a DISPLAY command, described later under "Displaying RMF Status."

Examples

Assume that RMF has been initialized with the system command:

```
START RMF.A,,,NOZZ
```

To start the Monitor I session, issue the command:

```
MODIFY RMF.A,START ZZ,NOCPU
```

As a result, the session is started, processor activity is not measured, and all other options are taken from other sources.

To modify the options to include measurement of processor activity, issue the command:

```
MODIFY RMF.A,MODIFY ZZ,CPU
```

To stop the Monitor I session from the operator console, issue the command:

```
MODIFY RMF.A,STOP ZZ
```


This command stops only the Monitor I session; other active sessions can continue processing, and RMF remains ready to accept further session commands.

Other examples of using the system MODIFY command to enter session commands that control session processing are given under "Monitor I Session Control," in Chapter 3, and under "Starting and Ending a Local 3270 Display Session" and "Monitor II Background Session Control" in Chapter 4.

Displaying RMF Status

Displaying RMF status enables the system operator to determine what sessions are active and what options are in effect for each active session. Note that TSO sessions are not included in the status information displayed because TSO sessions are completely controlled by the terminal user.

To display RMF status, issue a DISPLAY session command. The DISPLAY session command is passed to RMF through the input field of the system MODIFY command. The required syntax is:

```
                                         {ACTIVE     }
{MODIFY} {[RMF.]id,}  {DISPLAY}         {session-id }
{  F   } { jobname,}  {   D   }         {ALL        }
```

id
The identifier assigned to RMF when it was initialized.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

ACTIVE
Specifies that the session identifiers of all active non-TSO sessions are to be displayed. ACTIVE is the default value.

session-id
The session identifier for a particular session; when **session-id** is specified, the options for the named session are displayed.

ALL
Specifies that the session identifiers and current options for all active non-TSO sessions are to be displayed.

Examples

1. To display the session identifiers of all active non-TSO sessions, issue either the command:

   `MODIFY RMF.A,DISPLAY ACTIVE`

   or, because ACTIVE is the default, the command:

   `MODIFY RMF.A,DISPLAY`

2. To display the options for the Monitor I session, issue the command:

   `MODIFY RMF.A,DISPLAY ZZ`

3. To display the session identifiers and options for all active non-TSO sessions, issue the command:

   `MODIFY RMF.A,DISPLAY ALL`

RMF Termination

The STOP system command terminates RMF processing. Issuing this command causes all active non-TSO RMF sessions to end their processing. A message is issued informing you that RMF has terminated. The STOP and CANCEL system commands are the only ways to stop RMF once it has been initialized; methods of stopping individual sessions are described in Chapter 3 for the Monitor I session and in Chapter 4 for the Monitor II sessions. The syntax of the STOP system command is:

```
{STOP} {[RMF.]id}
{ P  } {jobname }
```

id
The identifier assigned to RMF when it was initialized by the START system command.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

Example

To terminate RMF processing (that is, to end the processing of all active non-TSO sessions), issue the command:

```
STOP RMF.A
```

RMF Options

While system, session, and display commands essentially control an RMF session, options tailor session processing to suit your particular needs. You specify options when the session is started and can also modify them during session processing. Seven parmlib members shipped with RMF contain possible options for running the various RMF session types. Chapter 3 gives a detailed description of RMF session options for Monitor I; chapter 4 gives a detailed description of session options for Monitor II.

RMF takes the options that control Monitor I sessions and Monitor II background sessions from the input sources in the following priority order:

1. START command parm value
2. EXEC statement PARM field (Monitor I session only)
3. One or more members of the RMF partitioned data set (usually SYS1.PARMLIB)
4. Program defaults

When the session starts, RMF merges the options from these sources in priority order. An option explicitly specified in the START command takes priority over any conflicting specifications in the EXEC statement or the RMF partitioned data set member or members. An option explicitly specified in the EXEC statement takes priority over a conflicting specification in the RMF partitioned data set member or members. If there are options for which no values are specified in the START command, the EXEC statement, or the RMF partitioned data set member or members, RMF uses the program defaults.

You can modify session processing during the session by specifying, on a MODIFY session command, the options to be changed. RMF changes only the options specified on the command. A changed option remains in effect until the session ends or the operator issues the MODIFY session command to change the option again.

Contents of the Parmlib Members

Seven parmlib members are shipped with RMF and are placed in SYS1.PARMLIB when RMF is installed. Four parmlib members apply to RMF sessions for Monitors I and II. Based on your installation's particular needs, you can modify options for Monitors I and II in these parmlib members or supply additional parmlib members. (For information about the parmlib members for a Monitor III session, see *MVS/XA Resource Measurement Facility (RMF) Monitor III Reference and User's Guide*, LC28-1557.) The parmlib members and their contents for Monitors I and II are:

- ERBRMF00 contains possible options for running the Monitor I session as a high-level measurement tool.
- ERBRMF01 contains possible options for a Monitor II background session.
- ERBRFM02 contains options that produce very detailed reports for a Monitor I session. This member is intended for use when you verify the installation of RMF.
- ERBRFM03 contains options that produce very detailed reports for a Monitor II background session. This member is intended for use when you verify the installation of RMF.

The options in each member are as follows:

ERBRMF00: The options supplied for the Monitor I session are:

```
CPU                   PAGESP
PAGING                NOTRACE
CHAN                  NOENQ
WKLD(PERIOD)          INTERVAL(60M)
DEVICE(NOTAPE)        NOSYNC
DEVICE(DASD)          NOSTOP
DEVICE(NOCHRDR)       CYCLE(1000)
DEVICE(NOUNITR)       NOEXITS
DEVICE(NOCOMM)        RECORD
DEVICE(NOGRAPH)       NOREPORT
DEVICE(NONMBR)        OPTIONS
DEVICE(NOSG)          SYSOUT(A)
IOQ(NOTAPE)           VSTOR(S)
IOQ(DASD)
IOQ(NOCHRDR)
IOQ(NOCOMM)
IOQ(NOUNITR)
IOQ(NONMBR)
IOQ(NOGRAPH)
```

ERBRMF01: The options supplied for a Monitor II background session are:

```
ASD              RECORD
NOUSER           REPORT(DEFER)
NODELTA          OPTIONS
SINTV(30S)       SYSOUT(A)
STOP(30M)
```

ERBRMF02: The installation verification options for the Monitor I session are:

```
CPU                        TRACE(RCVUICA,END)
PAGING                     TRACE(RCVCPUA,END)
CHAN                       TRACE(RCVASMQA,END)
WKLD(PERIOD,SYSTEM,        TRACE(RCVPTR,END)
     RANGE,DOMAIN,         TRACE(DMDTMPLT,END)
     OBJ,TIME)             TRACE(DMDTCMPL,END)
IOQ(DASD)                  TRACE(DMDTRUA,END)
IOQ(TAPE)                  INTERVAL(15M)
IOQ(CHRDR)                 SYNC(15M)
IOQ(UNITR)                 STOP(8H)
IOQ(COMM)                  CYCLE(250)
IOQ(NONMBR)                NOEXITS
IOQ(GRAPH)                 RECORD
DEVICE(TAPE)               REPORT(REALTIME)
DEVICE(DASD)               OPTIONS
DEVICE(CHRDR)              SYSOUT(A)
DEVICE(UNITR)              VSTOR(S)
DEVICE(COMM)
DEVICE(GRAPH)
DEVICE(NONMBR)
DEVICE(NOSG)
PAGESP
```

ERBRMF03: The installation verification options for a background Monitor II session are:

```
ASD              SENQR
ARD              TRX
ASRM             NOUSER
CHANNEL          DELTA
DDMN             SINTV(30S)
DEV              STOP(1H)
IOQUEUE          RECORD
PGSP             REPORT(DEFER)
SPAG             OPTIONS
SRCS             SYSOUT(A)
SENQ
```


Chapter 3. Monitor I Session

The measurement activities of a Monitor I session can be started at the operator console by a START command or as a batch job. See "Starting the Monitor I Session" on page 3-21. Only one Monitor I session can be active at a time.

The options that you select control session processing. Each option is described under "Monitor I Session Options." The options for a session can be specified in:

- The **parm** field of the START command that starts the session.
- The PARM field of the EXEC statement in the RMF cataloged procedure.
- The RMF partitioned data set member(s), normally included in SYS1.PARMLIB.
- The operator reply to message ERB306I.

The input from the sources is merged to form a final set of options for the session. When you do not specify an option in any of these sources, RMF uses a program default. Input merge is described under "Input Merge for the Monitor I Session" on page 3-17. The set of options established by input merge can be modified during the session, as described under "Modifying the Monitor I Session Options" on page 3-22.

Methods of ending the session are described under "Ending the Monitor I Session" on page 3-23. Examples and descriptions of Monitor I reports appear under "Monitor I Session Interval Reports" on page 3-25.

Monitor I Session Options

This section describes, in alphabetical order, the options that can be specified for a Monitor I session. The program defaults are underscored where appropriate.

<u>CHAN</u>/NOCHAN
: Specifies whether or not channel path activity is to be measured. A channel path is the physical interface that connects control units (grouped into logical control units) and devices to the channel subsystem.

<u>CPU</u>/NOCPU
: Specifies whether or not processor activity is to be measured.

  If you specify CPU **in a Processor Resource/Systems Manager (PR/SM) environment in LPAR mode,**, processor activity for all configured partitions is measured.

CYCLE(value)
: Specifies, in milliseconds, the length of the cycle at the end of which sampling observations are to be made. The **valid range** is from 50 to 9999 milliseconds. A specification of less than the minimum value (50 milliseconds) is increased to 50 milliseconds. A specification of more than the maximum value (9999 milliseconds) is decreased to 9999 milliseconds. The **default value** is 1000 milliseconds. See "INTERVAL and CYCLE Options" later in this section for considerations that apply to choosing a cycle length.

DEVICE (option,option...)/NODEVICE

Specifies whether or not system device activity is to be measured. You can request device activity by specifying all devices within one or more classes, or, optionally, one or more specific devices within each class.

Any of the following options can be specified:

- A device number in the form NMBR (nmbr1,nmbr2,...) where nmbr1 and nmbr2 are a one to three-digit hexadecimal number (leading zeroes can be omitted)
- NONMBR, which causes any lists of device numbers to be cancelled
- Any of the following classes:

CHRDR/NOCHRDR	-	Character reader devices
COMM/NOCOMM	-	Communications equipment
DASD/NODASD	-	Direct access storage devices
GRAPH/NOGRAPH	-	Graphics devices
TAPE/NOTAPE	-	Magnetic tape devices
UNITR/NOUNITR	-	Unit record devices

Note: Because of its speed, the IBM 3890 Document Processor is reported under unit record devices rather than character reader devices.

When you omit the device option, the defaults are as underscored above; that is, the device report is produced for direct access storage devices. If you explicitly specify DEVICE, you must include an option. However, the device option need include only the classes you want to add to the default (DASD) and/or the specific device number you want data for. For example, to request device reporting for magnetic tape devices and direct access storage devices, you would specify:

```
DEVICE(TAPE)
```

Direct access storage device reporting would be included by default, and the other device classes would, also by default, not be reported.

To limit the reporting of direct access storage devices, you must specify NODASD and use the NMBR field to identify those devices you want to monitor.

The NMBR field indicates that RMF is to report on the specific devices identified. The numbers can be expressed as a single device or as a range of devices. A range is indicated by specifying the first and last device numbers separated by a colon. Each single number or range is separated by a comma. For example, to request device reporting for magnetic tape devices 180, 183, 184, 185, and 188 as well as all direct access devices and communication equipment, you would specify:

```
DEVICE(COMM,NMBR(180,183:185,188))
```

Again, direct access device reporting would be included by default, and any device class (other than DASD or tape devices) would default to not being included in reporting.

NONMBR is the default; when specified, it causes any existing lists of device numbers to be cancelled.

Note: For more information on non-DASD/TAPE measurement, refer to the CMB parameter in IEASYS in *MVS/XA System Program Library Initialization and Tuning.*

To request device reporting for all storage groups, specify:

```
DEVICE(DASD,SG(aaa))
```

where aaa represents a storage group name. In the device report, RMF sorts the storage group names by national characters first, then alphabetically, and finally, by number. RMF always reports the storage group name of a volume when the volume is a member of a storage group, even if the SG suboption has not been selected. If the volume is added or deleted from a storage group, or if the storage management subsystem is not active, the storage group name may not be reported. If a volume does not belong to a storage group, the storage group field for that volume is blank, and it appears at the top of the report.

If you request the following:

```
DEVICE(NODASD, NMBR(288,291), SG(PROC01:PROC05))
```

the device report is divided into two parts. The first part of the report contains the devices specified by the NMBR suboption and is sorted by LCU and device number. The second part contains the devices specified for the SG suboption and is sorted by storage group and the device numbers within the group. Because you can specify a device on the NMBR suboption that is part of a storage group specified on the SG suboption, some devices might be reported twice.

NOSG is the default; when specified, it causes any existing lists of storage group names to be cancelled.

ENQ({DETAIL | SUMMARY} [,majorname[,minorname]])/NOENQ

Specifies whether or not enqueue contention activity is to be measured. When DETAIL is specified, the report includes the names of one or two jobs owning the resource and one or two jobs that have the longest period of contention and are waiting for the resource. The names reported are selected during the period of maximum contention. No names are reported when you specify SUMMARY. When you omit the ENQ option, the default is NOENQ. If you explicitly specify ENQ, you must specify either DETAIL or SUMMARY; specifying ENQ without choosing DETAIL or SUMMARY causes a syntax error.

The optional **majorname** field can contain the one to eight character major name of a serially-reusable resource. Optionally, the major name is followed by a comma and a minor name. The **minorname** field can contain the one to forty-four character minor name of the resource. For example:

```
ENQ(DETAIL,MAJORNAME,MINORNAME)
```

You can use the name fields to measure the contention for a certain specific resource; when you omit a specific name, all resources for which contention occurs are measured. When you omit a minor name, all resources with the major name you specify are included.

EXITS/NOEXITS

Specifies whether or not Monitor I user exit routines are to be executed during session processing to enable your installation to gather and report data in addition to the data measured by RMF. See "Monitor I Session User Reports" on page 6-2 for information on using the exit routines.

INTERVAL $\left\{ \begin{array}{c} \underline{\textbf{30M}} \\ \textbf{value[M]} \end{array} \right\}$

Specifies, in minutes (M), the length of the RMF reporting interval. At the end of each interval, RMF is dispatched. All data is gathered and sampled data is summarized. The data is then formatted for an interval report or an SMF record. Either or both can be done. SMF records are written at this point, while reports can be written or held for later printing, depending on the value you specify for the REPORT option. The **valid range** for INTERVAL is from 1 to 60 minutes. Specifying a value outside this range causes RMF to use 1 minute if the value is below the range and 60 minutes if the value is above the range. The **default range** is thirty minutes (30M). You can synchronize the RMF reporting interval to any time within the hour by using the SYNC option. See the SYNC/NOSYNC option later in this section for more information.

When RMF is executing and does not get control within 99 minutes, no data is collected. RMF issues a message to tell the operator that data collection was skipped for the interval. This problem can occur when the dispatching priority for RMF is too low; see "RMF JCL" in Chapter 2 for information about changing the dispatching priority. The interval length can also exceed 99 minutes when the processor is stopped during the interval; to avoid skipped data collection in this case, terminate the session or RMF before stopping the processor.

When the IPS or installation control specification changes during an interval when RMF is measuring workload activity, RMF performs normal end-of-interval processing when the change occurs. That is, RMF generates a set of SMF records and, if the REPORT option is in effect, reports for the part of the interval that elapsed before the change. RMF then reinitializes all active measurements and continues normal processing using the new values for the IPS or installation control specification. Although RMF automatically stops and restarts the session, you should be aware of these actions because no data is sampled during the brief period required to stop, modify, and restart the session.

IOQ (option,option...)/NOIOQ

Specifies whether or not I/O queuing activity is to be measured for the devices in each logical control unit. The definition of a logical control unit (LCU) is model-dependent.

On a 308x or 908x processor, an LCU is the set of devices attached to the same physical control unit (or group of control units that have one or more devices in common.) Each device belongs to only one LCU and all requests to the various devices within an LCU share the same LCU queue for I/O request initiation and redrive within the channel subsystem.

On a 4381 processor complex, an LCU is the set of devices attached to an identical set of physical control units. A physical control unit can belong to more than one LCU if the devices are configured asymmetrically. Each device belongs to only one LCU. All I/O requests and redrives are queued on a physical control unit basis.

On a 3090 processor, an LCU is the set of devices attached to the same physical control unit (or group of control units that have one or more devices in common). As with the 308x processor, each device belongs to only one LCU, but the I/O processor (IOP), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU of a 3090 processor.

On all processors, you can request I/O queuing activity by specifying all LCUs within one or more classes, or, optionally, one or more specific LCUs within each class.

Any of the following options can be specified:

- An LCU number in the form NMBR (LCUN1,LCUN2,...) where LCUN1 and LCUN2 are a one to three-digit hexadecimal number or NMBR(LCUN1:LCUN2,...) where LCUN1 and LCUN2 are one to three-digit hexadecimal numbers that indicate a specific range of LCU numbers.
- NONMBR, which causes any lists of LCU numbers to be cancelled
- Any of the following classes:

CHRDR/NOCHRDR	-	Character reader
COMM/NOCOMM	-	Communications equipment
DASD/NODASD	-	Direct access storage
GRAPH/NOGRAPH	-	Graphics
TAPE/NOTAPE	-	Magnetic tape
UNITR/NOUNITR	-	Unit record

When you omit the IOQ option, the defaults are as underscored in the preceding list; RMF measures I/O queuing for all logical control units (LCUs) of the DASD device class. If you explicitly specify IOQUEUE, you must include an option; however, the option need include only the classes you want to add to the default (DASD) and/or the specific LCU number you want data for.

For example, to request I/O queuing activity for LCUs for magnetic tape devices and direct access storage devices, you would specify:

```
IOQ(TAPE)
```

Logical control units of the direct access storage devices would be included in the report by default, and the other device classes would also by default not be reported.

To limit the reporting to only some LCUs for direct access storage devices, you must specify NODASD and use the NMBR field to identify those LCUs you want to monitor.

The NMBR field indicates that RMF is to provide I/O queuing activity for the specific LCU numbers identified. The numbers can be expressed as a single LCU number or as a range of LCU numbers. A range is indicated by specifying the first and last LCU numbers separated by a colon. Each single LCU number or range is separated by a comma. For example, to request I/O queuing activity for magnetic tape device logical control units 130, 133, 134, 135, and 150 as well as all LCUs of the DASD and COMM classes, you would specify:

```
IOQ(COMM,NMBR(130,133:135,150))
```

Again, LCUs of direct access storage devices would be included by default, and the other device classes would also by default not be reported.

NONMBR is the default; when specified, it causes any existing lists of LCU numbers to be cancelled.

MEMBER(list)
Specifies the member or members of a partitioned data set that contain the options to be used for the session. The list can contain from one to five entries, separated by commas, and each entry consists of a two-character alphameric

value to be appended to ERBRMF to form the member name. When more than one member is specified, input merge takes the options from the members in left to right priority order.

For a Monitor I session, the default is 00, indicating member ERBRMF00 in the partitioned data set named on the IEFRDER DD statement in the RMF cataloged procedure (normally SYS1.PARMLIB). The MEMBER option must not be specified within a partitioned data set member.

For a list of the options specified in the members supplied by IBM, see "Contents of the Parmlib Members" in Chapter 2.

Each member specified must contain options appropriate for the Monitor I session. A member containing Monitor II background session options will cause syntax errors.

NOZZ

Specifies that the Monitor I session is not to be started when RMF is initialized. NOZZ indicates RMF initialization only. NOZZ is mutually exclusive with all other options; that is, when NOZZ is specified, it must be the only option specified. When NOZZ is specified on the EXEC statement and RMF is initialized by a START command that does not specify NOZZ, the NOZZ option on the EXEC statement is ignored, and RMF starts the Monitor I session as specified on the START command.

NOZZ can be specified only on the START command for RMF or in the PARM field of the EXEC statement. It must not appear within a partitioned data set member or in an operator reply to an RMF message.

OPTIONS/NOOPTIONS

Specifies whether or not a list of the options to be used is to be printed at the operator console at the start of the session. If the list is printed (OPTIONS specified), the operator can respond with any desired changes, except the MEMBER option.

On the operator command, OPTIONS can be abbreviated OPTN, and NOOPTIONS can be abbreviated NOOPTN.

You can avoid unnecessary console output and delay in starting the session by specifying NOOPTIONS. However, if RMF detects any syntax errors during input merge, OPTIONS is forced.

Figure 3-1 shows an example of the console output produced when OPTIONS is in effect. Each option is followed by the input source from which RMF input merge obtained the option. The possible sources are:

- COMMAND -- the option was specified on a START or MODIFY command.
- DEFAULT -- the option was taken from the program defaults.
- EXEC -- the option was specified on the EXEC statement in the RMF cataloged procedure.
- CHANGED -- the option was changed by RMF. A preceding message will describe the nature of the conflict and the change RMF has made.
- MEMBER -- the option was specified in a member of the RMF partitioned data set.
- REPLY -- the option was changed by the operator in reply to message ERB306I.


```
ERB305I ZZ : PARAMETERS
ERB305I ZZ :  WKLD (PERIOD) -- DEFAULT
ERB305I ZZ :  NOSYNC -- DEFAULT
ERB305I ZZ :  TRACE (RCVPTR,ALL) -- MEMBER
ERB305I ZZ :  TRACE (RCVASMQA,ALL) -- MEMBER
ERB305I ZZ :  TRACE (RCVCPUA,ALL) -- MEMBER
ERB305I ZZ :  TRACE (RCVUICA,ALL) -- MEMBER
ERB305I ZZ :  NOEXITS -- MEMBER
ERB305I ZZ :  ENQ(DETAIL) -- MEMBER
ERB305I ZZ :  OPTIONS -- COMMAND
ERB305I ZZ :  STOP(168H) -- MEMBER
ERB305I ZZ :  CYCLE(250) -- MEMBER
ERB305I ZZ :  INTERVAL(15M) -- MEMBER
ERB305I ZZ :  RECORD -- MEMBER
ERB305I ZZ :  SYSOUT(z) -- MEMBER
ERB305I ZZ :  REPORT(REALTIME) -- MEMBER
ERB305I ZZ :  PAGESP -- MEMBER
ERB305I ZZ :  DEVICE(NOCHRDR) -- DEFAULT
ERB305I ZZ :  DEVICE(NONMBR) -- DEFAULT
ERB305I ZZ :  DEVICE(GRAPH) -- MEMBER
ERB305I ZZ :  DEVICE(COMM) -- MEMBER
ERB305I ZZ :  DEVICE(DASD) -- MEMBER
ERB305I ZZ :  DEVICE(TAPE) -- MEMBER
ERB305I ZZ :  DEVICE(UNITR) -- MEMBER
ERB305I ZZ :  DEVICE(NOSG) -- DEFAULT
ERB305I ZZ :  CHAN -- MEMBER
ERB305I ZZ :  PAGING -- MEMBER
ERB305I ZZ :  IOQ(NOCHRDR) -- DEFAULT
ERB305I ZZ :  IOQ(NONMBR) -- MEMBER
ERB305I ZZ :  IOQ(GRAPH) -- MEMBER
ERB305I ZZ :  IOQ(COMM) -- MEMBER
ERB305I ZZ :  IOQ(DASD) -- MEMBER
ERB305I ZZ :  IOQ(TAPE) -- MEMBER
ERB305I ZZ :  IOQ(UNITR) -- MEMBER
ERB305I ZZ :  VSTOR(S) -- MEMBER
ERB305I ZZ :  MEMBER(04) -- COMMAND
```

Figure 3-1. Console Output From the OPTIONS Option

<u>PAGESP</u>/NOPAGESP
Specifies whether or not page/swap data set activity is to be measured.

<u>PAGING</u>/NOPAGING
Specifies whether or not system paging activity is to be measured.

<u>RECORD</u>/NORECORD
Specifies whether or not measured data is to be written to the SMF data set. In order for RECORD to take effect, the complementary SMF enabling procedures must first be performed. These procedures are described in *System Management Facilities (SMF)*.

REPORT({REALTIME | DEFER})/<u>NOREPORT</u>
Specifies whether or not printed interval reports of the measured data are to be produced. When you omit the option, the default is NOREPORT. If you explicitly specify REPORT, you must specify either REALTIME or DEFER; specifying REPORT without choosing either REALTIME or DEFER causes a syntax error.

REALTIME indicates that the reports are to be printed when formatted at the conclusion of the interval; DEFER indicates that the reports are to be printed after RMF processing terminates.

<u>STOP</u> (value $\begin{bmatrix} \underline{\mathbf{M}} \\ \mathbf{H} \end{bmatrix}$ **)/NOSTOP**

Specifies the desired duration for the Monitor I session in minutes (M) or hours (H). The **valid range** is from one minute to one week (168 hours or 10,080 minutes). If you do not specify a value, the **default range** is 8 hours. If you specify a value outside the range, RMF uses one minute if the value is below the range and 168 hours if the value is above the range. If neither M nor H is specified, M (minutes) is assumed. NOSTOP means that the session can be ended only by a STOP command. Note that the STOP option applies only to the session. RMF remains active until the operator issues a STOP system command.

The operator STOP command can end all the sessions at any time, regardless of the value specified for this option, provided that a session identifier was specified or assigned automatically when the session was started.

Because of SYSOUT space limitations, STOP (interval) will be forced when both NOSTOP and REPORT(DEFER) are specified, where **interval** is the value of the INTERVAL option after it has been validated during input merge.

SYNC (nnM)/<u>NOSYNC</u>

Specifies whether or not the interval is to be synchronized with the hour. The **valid range** is the number of minutes from 0 to 59 (nn), after the hour at which synchronization is to occur. If any value other than 0 through 59 is specified, RMF assigns a **default value** of 0. RMF synchronizes the interval by shortening the first interval. Subsequent intervals remain synchronized only when the length of the specified interval can be evenly divided into 60. For example, if you specify an interval of 20 minutes synchronized on 10 minutes, reports are generated at 10, 30, and 50 minutes past the hour. Therefore, if you start your session at 9:05, the first interval is shortened so that a report is generated at 9:10. Similarly, if you start your session at 9:15, the first interval is shortened so that a report is generated at 9:30.

NOSYNC is the default and specifies that intervals are not to be synchronized.

SYSOUT(class)

Specifies the SYSOUT class to which the formatted interval reports are directed. Class A is the default. The SYSOUT option cannot be modified during the session.

TRACE(fieldname[,opt list])/<u>NOTRACE</u>

Specifies whether or not the contents of ASM/RSM/SRM control block fields, or certain SMF data fields are to be traced, where:

fieldname

Is a valid field name in one of the following:

- Auxiliary storage manager vector table (ASMVT)
- System resource management (SRM) data areas
- System resource management (SRM) domain table
- RSM control and enumeration area (RCE)
- System management facilities (SMF) data area (SMCA)
- A user-specified field name

Note: For information on how to enable RMF to trace the contents of a user-specified field, see "Tracing Your Own Field" in Chapter 7.

The RCE fields that can be specified and their meanings are:

field	**meaning**
RCEAFC	total number of frames currently on all available frame queues
RCEPOOL	number of frames currently available to the system, including frames backing permanent storage (nucleus frames, hardware storage area frames, FLPA frames or fixed BLDL frames), bad frames and offline frames.
RCEDFRS	number of times a deferred frame allocation has been satisfied
RCETOTPI	total number of pages paged-in excluding swap-in and VIO page-in
RCETOTPO	total number of pages paged-out, excluding swap-out, move-out of VIO pages, and page-out of VIO pages
RCEESREA	number of pages read from extended storage
RCEVIOME	number of VIO pages written to extended storage
RCEVIORE	number of VIO reads from extended storage
RCEVIOMG	number of VIO pages migrated from extended storage to paging data sets
RCEVIOPI	total number of VIO pages paged-in, excluding swap-in
RCEVIOPO	total number of VIO pages, excluding swap-out, moved out or paged-out
RCEVIORU	number of times a VIO data set page was reused
RCESWPPI	total number of pages requiring I/O to swap-in
RCESWPPO	total number of pages requiring I/O to swap-out
RCETOTRC	total number of times a page was reclaimed from an available frame queue
RCECOMPI	number of common area pages paged-in
RCECOMPO	number of common area pages paged-out
RCECOMRC	number of common area pages that have been reclaimed from an available frame queue
RCECOMAL	number of frames allocated to common
RCESPFR	number of frames available by swap-out without requiring I/O
RCETOTFX	total number of pages currently fixed, the sum of page-fixed LSQA, SQA (excluding reserved SQA) and V=R allocated pages
RCELPAPI	number of PLPA and PLPA directory pages paged-in
RCELPARC	number of PLPA and PLPA directory pages that have been reclaimed from an available frame queue
RCEPAGMV	number of times a frame was moved from one frame to another
RCEESINU	number of in-use extended storage frames (3090 processor only)
RCEESWRT	number of pages written to extended storage frames (3090 processor only)
RCEMVBEL	number of pages moved out of real storage below the 16 megabyte line
RCENWSF	total number of secondary and non-working set pages migrated to
RCEWSDNE	total number of primary working set pages migrated to auxiliary storage (3090 processor only)

The ASMVT fields that can be specified and their meanings are:

field	**meaning**
ASMSLOTS	total local slots (sum of slots in open local page data sets)
ASMVSC	total local slots allocated for VIO private area pages
ASMNVSC	total local slots allocated for non-VIO private area pages
ASMERRS	bad slots on local page data sets
ASMIORQR	count of I/O requests received by I/O control
ASMIORQC	count of I/O requests completed and returned to RSM

The SRM data area fields that can be specified and their meanings are:

field	meaning
RMCAUSSC	unilateral swap out count
RMCANQSC	CPU enqueue exchange swap count
RMCATOSC	terminal output swap count
RMCAEXSC	exchange on recommendation value swap count
RMCARSSC	real storage shortage swap count
RMCATISC	terminal input swap count
RMCALWSC	long wait physical swap count
RMCAXSSC	auxiliary storage shortage swap count
RMCADWSC	detected wait physical swap count
RMCARQSC	requested swap count
RMCAFHLD	number of swaps failed because of an outstanding HOLD SYSEVENT
RMCATSSC	count of transition swaps
CCVRBSWT	recent base system wait time
CCVRBSTD	recent base time of day
CCVEJST	low priority user job step time
CCVIMBBT	time of last CPU imbalance
CCVUTILP	system CPU utilization
CCVLGUTL	long term CPU utilization * 256
CCVENQCT	number of users non-swappable for enqueue reasons
CCVCPUCT	number of online CPUs
MCVSTCRI	highest system UIC
MCVFRCNT	number of pages needed to be stolen by force steal routine
MCVMGAGE	extended storage migration rate (3090 processor only)
RCVUICA	UIC average
RCVCPUA	CPU usage average * 16
RCVAVQC	AVQ low count
RCVASMQA	ASM queue length average
RCVPTR	paging rate
RCVDPR	demand paging rate
RCVPAGRT	total paging rate
RCVSWPTM	time (in milliseconds) used by ASM to process a request to transfer a group of pages to or from a data set
RCVMFXA	average number of fixed frames for the system
RCVFXIOP	percentage of real storage that is fixed or allocated for paging
RCVMSPP	page delay time excluding swap paging (in milliseconds)
LSCTCNT	current number of logically swapped users for terminal wait
LSCTMTE	maximum think time allowed for logical swap candidate
MCVSIPR	common page-in rate
MCVTWSS	common target working set size
RMCTTRPC	number of pages used for transaction elements
MCVSBLTF	long term percentage of eligible storage that is actually fixed
MCVSBLTS	long term steal criteria average
MCVSBFQA	short term available frame count average
MCVSBRVF	storage contention factor used for computing the storage load balancing recommendations
MCVSBIMT	time of day when storage was last considered out of balance
RCVAFQA	average available frame count
MCVSBSIG	frame count threshold for significant storage users

The domain table fields that can be traced and their meanings are:

field	meaning
DMDTMPLT	current multiprogramming level target
DMDTCMPL	current multiprogramming level
DMDTINCU	users per domain in storage
DMDTOUTU	users per domain out of storage
DMDTRUA	average number of ready users * 16
DMDTTWSR	time weighted service rate for a domain
DMDTCIDX	domain contention index * 256
DMDTTWET	average elapsed time in 1024 microsecond units for first period TSO transaction (does not include delay imposed by SRM to meet response time objective)

The domain fields appear in the report as follows:

field specified	field reported
DMDTINCU	DMINUxxx
DMDTOUTU	DMOUTxxx
DMDTMPLT	DMTARxxx
DMDTRUA	DMRUAxxx
DMDTCMPL	DMMPLxxx
DMDTTWSR	DMTWSxxx
DMDTCIDX	DMCDXxxx
DMDTTWET	DMTWExxx

The SMF data fields that can be traced and their meanings are:

field	meaning
SMCANMFL	current number of full buffers
SMCACNBF	current number of buffers
SMCABFWT	number of buffers written
SMCARCWT	number of records written
SMCABFLS	number of records lost because of a shortage of buffers
SMCADSCT	number of records lost because of a full data set

The logical path block (LPB) field that can be traced and its meaning is:

field	meaning
LPBUTIL	LPB utilization

opt list

Is one or more of the following:

MIN	the minimum sampled value of the field over the sampling period.
MAX	the maximum sampled value of the field over the sampling period.
AVG	the average value of the field over the sampling period.
END	a snapshot of the last value in the sampling period.
STDDEV	the standard deviation from the values sampled.
<u>ALL</u>	the default for opt list, meaning all of the above.

VSTOR[({S | D} [,jobname1,jobname2,...])] /NOVSTOR

Specifies whether or not virtual storage activity is to be measured. RMF can produce common storage summary and detail reports and private area summary and detail reports. When you specify S, either explicitly or by default, RMF produces summary reports; when you specify D, RMF produces both summary reports and detail reports. (Specifying S or D affects only the reports RMF produces; RMF always collects the data required for a detail report.)

To obtain private area reports, replace **jobname** with the name of the job to be reported. RMF gathers private area data only when you specify a job name. While the syntax allows you to specify the names of up to 25 jobs, it is more practical to minimize the time required to gather the data by specifying one or two jobs at any one time. When selecting specific jobs, note also that RMF can gather meaningful data only for long-running jobs.

If you omit the VSTOR option, the default is VSTOR(S). If you specify VSTOR without any operands, RMF also produces a summary report for common storage. Examples of other possible combinations are:

- VSTOR(D) produces a summary and detail report for common storage.
- VSTOR(D,VTAM) produces a summary and detail report for common storage and a summary and detail report for the private area of the VTAM address space.
- VSTOR(MYJOB) produces a summary report for common storage and a summary report for the private area of the MYJOB address space.

If you specify the name of a job that is not running when RMF begins measuring virtual storage activity, RMF issues a message indicating that it cannot gather data about the named job. For as long as the VSTOR option remains unchanged, RMF searches for the job at the beginning of each interval. The message appears on the operator console and in the SYSOUT message data set; when RMF finds the job, it deletes the message from the operator console.

WKLD(list)/NOWKLD

Specifies whether or not system workload activity is to be measured. When WKLD is specified, **list** must indicate the type of reporting required. Any or all of the following can be specified in **(list)**:

PERIOD/NOPERIOD	--	Requests reporting by performance group period type 1 sub-report)
GROUP/NOGROUP	--	Requests reporting by performance group (type 2 sub-report)
RANGE/NORANGE	--	Requests reporting by performance objective, domain, and performance group in a range of performance group numbers (type 3, 4, and 5 sub-reports)
OBJ/NOOBJ	--	Requests reporting by performance objective (type 6 sub-report)
DOMAIN/NODOMAIN	--	Requests reporting by domain number (type 7 sub-report)
SYSTEM/NOSYSTEM	--	Requests reporting by system (type 8 sub-report)
TIME/NOTIME	--	Requests reporting by time slice group (type 9 sub-report)

When you omit the option, the defaults are as underscored above. When you explicitly specify WKLD, you must include **list** or a syntax error occurs.

Need for Careful Choice of Options

Care should be taken in specifying certain groups of RMF options. Some options cannot be used concurrently to control a single Monitor I session. Should any conflicts occur, RMF detects the mutually-exclusive options during input merge and selects compatible values for these options; the operator is notified of the selections made. The possible conflicts are described in Figure 3-2.

Figure 3-2. RMF Response to Conflicting Options

Conflict	Problem	RMF Resolution
NOREPORT AND NORECORD specified	No way for installation to obtain measurement data	Change NOREPORT to REPORT (DEFER)
STOP value specified is less than INTERVAL	Indicates session termination before obtaining any data	Set STOP value equal to INTERVAL value
REPORT(DEFER) and NOSTOP specified	SYSOUT becomes cluttered with unprinted reports	Change NOSTOP to STOP value equal to INTERVAL value

Other groups of options do not cause actual conflicts, but their values must be chosen carefully to avoid undesirable results. Those options that must be specified with care include:

- INTERVAL values and CYCLE values
- STOP, INTERVAL, and REPORT values
- RECORD, TRACE, INTERVAL, and CYCLE values
- Device class selection for the DEVICE option

INTERVAL and CYCLE Options

Much of the data in the paging, page/swap data set, processor, ASM/RSM/SRM trace, virtual storage, CPU, I/O queuing, and device activity reports is statistically sampled. Because, according to statistical theory, the accuracy of sampled data increases with the number of samples taken of random events, you would expect to observe more precise results with decreased CYCLE time (for a fixed INTERVAL value), or with increased INTERVAL length (for a fixed CYCLE value). For example, 400 samples taken of random independent events provide a value that, with 90% confidence, should fall within 4% of the true value; 1,600 samples of random independent events decrease to 2% the expected range of error, with 90% confidence.

However, pure statistical predictions are not always applicable to a software measurement tool such as RMF because the assumptions on which they are based (unbiased random independent samples and an infinite population) might not hold in an operating environment. Bias might occur because RMF samples internal indications of external system events. Thus, RMF values might not precisely approach the values measured by a hardware measurement tool.

The independence assumption becomes less and less realistic as CYCLE gets very small. As CYCLE gets smaller, each sample is more likely to find the system performing the same functions as in the previous sample; therefore, the new sample adds little additional information. The use of a smaller CYCLE value (while holding INTERVAL constant) should not be detrimental to accuracy, but any increase in accuracy might be of questionable benefit when compared with the system overhead that is introduced. A reasonable minimum CYCLE value is a function of the timing characteristics of the hardware being measured.

Note: In a 3090 processor, RMF acquires data about the use of channels through the system activity display (SAD). The hardware maintains running counters and samples at a frequency far exceeding that of any software mechanism. As a result, the frequency at which RMF picks up the counts on channel use does not affect the accuracy of the measurements. Other values acquired on a 3090 processor, such as MIN and MAX values for the available extended storage frames, do depend on sampling rather than on running counters. The sampling rate, as described earlier, does affect the expected accuracy of the information recorded in these fields.

STOP, INTERVAL, and REPORT Options

As mentioned earlier, the specification of NOSTOP along with REPORT(DEFER) is considered a conflict by RMF, because of the possible filling up of SYSOUT spool space. A similar problem can occur when the STOP value specified is very large, the INTERVAL value is small, and REPORT(DEFER) is specified.

RECORD, TRACE, INTERVAL, and CYCLE Options

When writing SMF trace records (RECORD specified) use discretion when specifying INTERVAL, CYCLE, and the desired trace options. SMF restricts the size of records written by the user to 32K bytes; therefore, the SMF record (type 76) will be truncated after the last complete set that will fit within 32,756 bytes. However, the truncated "set data" is still included in the fields of the trace record that contain totals for all set data.

Note: Trace reports printed with the REPORT option are not affected by a truncated SMF record and will contain all the set data. However, trace reports printed by the post processing routine from the same record would lack the truncated set data.

The following formula should be used in calculating the type 76 record size:

$$\text{record size} = 104 + \left(\frac{\text{interval}}{\text{cycle} * \text{samples per set}}\right) * \text{set size}$$

104
: Indicates the length, in bytes, of the SMF record common header, header extension and RMF product section.

interval
: Indicates the value specified on the INTERVAL option, converted to milliseconds.

cycle
: Indicates the value specified on the CYCLE option in milliseconds.

samples per set
: Indicates the IBM-supplied default of 60 unless the system programmer uses the SUPERZAP macro. (See "ASM/RSM/SRM Trace Activity Report" later in this chapter.)

set size
Indicates the length, in bytes, of the array that contains the set data for the field name. This value depends on the length of the data field traced and the trace options specified as follows:

Trace Option	4-byte data field	2-byte data field
MIN	4	2
MAX	4	2
AVG	8	6
END	4	2
STDDEV	20	14
STDDEV and AVG	20	14
ALL	32	20

Note: If the IBM-supplied default of 60 samples per set is used with a cycle value equal to or greater than 60 milliseconds, data will never be truncated from the SMF record.

Examples of Calculating Record Sizes

1. If the interval value specified is 30 minutes (1,800,000 milliseconds), the cycle value specified is 250 milliseconds, and the trace options specified for RCEPOOL (a 2-byte data field) are minimum, maximum, and standard deviation, the resulting calculation would be:

   $$\text{record size} = 104 + \left(\frac{1{,}800{,}000}{250 * 60}\right) * 18 \text{ or } 2{,}264 \text{ bytes}$$

   Notes:

   - The IBM-supplied default of 60 was in effect for samples per set.
   - MIN for a 2-byte field name requires 2 bytes, MAX requires 2 bytes, and STDDEV requires 14 bytes, totaling 18 bytes for the set size value.

2. If the interval value specified is 60 minutes (3,600,000 milliseconds), the cycle value specified is 100 milliseconds, all the trace options for ASMSLOTS (a 4-byte data field) are specified, and the samples per set value was zapped to 30, the resulting calculation would be:

   $$\text{record size} = 104 + \left(\frac{3{,}600{,}000}{100 * 30}\right) * 32 \text{ or } 38{,}504 \text{ bytes}$$

   A record size of 38,504 causes the last 174 sets to be truncated.

3. If the interval value specified is 60 minutes (3,600,000 milliseconds), the cycle value specified is 60 milliseconds, all the trace options for RMCAUSSC (a 4-byte data field) are specified, and the samples per set value is the IBM-supplied default of 60, the resulting calculation would be:

   $$\text{record size} = 104 + \left(\frac{3{,}600{,}000}{60 * 60}\right) * 32 \text{ or } 32{,}104 \text{ bytes}$$

Device Class Selection for the DEVICE Option

Because RMF overhead is directly related to the number of devices being measured, the DEVICE option list should include only those devices that require measurement. To further reduce RMF overhead, select specific devices for reporting rather than entire device classes. A good example of a candidate for reporting by device number to reduce system overhead would be the mass storage system (MSS). In addition to reducing RMF overhead, limiting the devices can also decrease the amount of SQA storage required by a Monitor I session as well as decrease the length of the SMF device records. In the case of post processor routines, selecting specific devices can result in shorter reports, thus saving both time and paper. Your storage administrator uses the storage management subsystem to define storage groups. Storage groups are a set of DASD volumes that have been assigned one common name. By using storage groups, volumes can be grouped together in easily measurable sets. For example, assign storage groups with paging volumes separate from storage groups with excessively used data sets.

The values you specify for the CYCLE option and the interval option also affect overhead. By decreasing CYCLE length or increasing INTERVAL length, you can increase sample size (number of samples per interval). Note, however, that decreasing the CYCLE length could significantly degrade system performance, especially in the device measurements area. At each sampling cycle, there is system overhead introduced by processing the periodic interruption and collecting the data. Because device data collection is the major contributor to this type of overhead, installations with a large number of UCBs (internal control blocks describing the devices generated with the system) can use the following approximation to estimate the total overhead (number of instructions executed) for each sampling cycle:

```
# of instructions = 70 + (10 * D1) + (60 * D2) + (100 * D3)
```

D1
Indicates the number of UCBs associated with offline devices for monitored device classes.

D2
Indicates the number of UCBs associated with online devices for monitored device classes.

D3
Indicates the average number of UCBs that have an I/O request outstanding on each sample. The number might be about 5% of the number of active devices, but can vary greatly depending on system workload.

Note: When a condition (such as device allocated, device reserved, device not ready, and mount pending) is tested and found to be true for an online device, the number of instructions executed for each sampling cycle increases by approximately 5 for each true condition. If the device is a multiple exposure device (such as an IBM 2305 Fixed Head Storage Device), the instruction count increases by approximately 17.

Therefore, the cycle value should not be made too small, especially when the number of UCBs for measured device classes is large.

Input Merge for the Monitor I Session

Input merge is the process RMF performs to determine what options are to be established for the Monitor I session. Input merge is performed at the start of the session and whenever the MODIFY session command is used to modify the options of an active Monitor I session. Input merge takes options from various input sources and merges them to form a list of options for the session.

When the Session Is Started

At the start of the session, input merge uses a "fill in the blanks" procedure; that is, it takes options from the four possible input sources in a defined priority order. Once an option is filled in, no options from a lower-priority input source can override that option. The input sources and their priority order are:

1. The **parm** field of the START command. Any options specified in this input source have the highest priority. If the MEMBER option is specified, the options contained in the member(s) are processed after any other options specified on the START command and before RMF proceeds to the next input source.

   The MEMBER option can be used to specify up to five members of the RMF partitioned data set (normally SYS1.PARMLIB). If the MEMBER option is used to specify more than one member, RMF takes options from the members in left to right order. If, for example, MEMBER(02,07) is specified, RMF generates the member names ERBRMF02 and ERBRMF07 and takes options from ERBRMF02 first. Thus, if ERBRMF02 specifies ENQ(DETAIL) and ERBRMF07 specifies ENQ(SUMMARY), ENQ(DETAIL) is the option established for the session. If a specified member does not exist, RMF continues with the next specified member, if any, or proceeds to the next input source.

2. The PARM field of the EXEC statement in the RMF cataloged procedure. Any options specified in this input source have the second highest priority. Options specified here cannot override any options from the **parm** field of the START command. If, for example, CHAN is specified on the START command and NOCHAN is specified on the EXEC statement, CHAN is established during input merge.

   If the MEMBER option is specified, the options contained in the member(s) are processed after any options specified on the EXEC statement and before RMF proceeds to the next input source. Use of the MEMBER option is as described earlier for item 1. Note that the EXEC statement is not used when the Monitor I session is started by a START session command.

3. If the MEMBER option has not been specified on either the START command or the EXEC statement, RMF uses a default member name, ERBRMF00, as the third highest priority input source. Options from the default member can fill in only those options not specified earlier on a higher priority input source.

4. The program defaults. The program defaults are the lowest priority input source; that is, the program defaults can fill in only those options not filled in by a higher-priority input source.

During input merge, RMF might detect mutually exclusive options or syntax errors. When either occurs, RMF notes the condition and issues a warning message, but it does not terminate the merge. When RMF detects mutually exclusive options (such as both RECORD and NORECORD), it takes the value associated with the first specification for the option. When an invalid value is detected for CYCLE, INTERVAL, SYNC, or STOP, RMF substitutes a valid value, as described for each option under "Monitor I Session Options." When other errors are detected, RMF ignores the input option in error and uses the next valid specification from a lower priority source.

At the end of input merge, RMF checks for conflicting options; see Figure 3-2 for a description of these options and the actions RMF takes. RMF displays the options resulting from input merge at the operator console when:

- OPTIONS is in effect.
- An error has been detected during the merge.
- Conflicting options have been detected and changed.

The system operator can change the options, if necessary, before session processing begins.

Examples of the Input Merge Process

To show how the input merge process works, assume:

1. RMF partitioned data set member ERBRMF02 contains the options:

   ```
   WKLD(SYSTEM)
   OPTIONS
   NOPAGESP
   EXITS
   ```

2. RMF partitioned data set member ERBRMF10 contains the options:

   ```
   NOEXITS
   DEVICE(NOUNITR,TAPE)
   ```

3. The RMF cataloged procedure contains the following statement:

   ```
   //IEFPROC EXEC PGM=ERBMFMFC,REGION=4096K,
   //             PARM='CYCLE(150),DEVICE(NOTAPE,DASD),
   //             MEMBER(02)'
   ```

If the operator initializes RMF and starts the Monitor I session with the following START command:

```
START RMF.A,,,(WKLD(GROUP),MEMBER(10),CYCLE(100),
 DEVICE(COMM))
```

Input merge uses the three options from the START command to begin the list of session options:

```
WKLD(GROUP)
CYCLE(100)
DEVICE(COMM)
```

Next, RMF generates the member name ERBRMF10 from the MEMBER option specified on the START command, and takes options from the member to continue filling in the list of options. After processing ERBRMF10, the option list consists of:

```
WKLD(GROUP)
CYCLE(100)
DEVICE(COMM,NOUNITR,TAPE)
NOEXITS
```

Input merge then proceeds to the second highest priority source, the PARM field of the EXEC statement in the RMF cataloged procedure. CYCLE(150) and

DEVICE(NOTAPE) are ignored because these options have been filled in by a higher-priority source. RMF adds:

```
DEVICE(DASD)
```

to the list of options and generates the member name ERBRMF02 from the MEMBER option specified on the EXEC statement. WKLD(SYSTEM) and EXITS in member ERBRMF02 are ignored because these options have been filled in. After processing the other options contained in ERBRMF02, the option list consists of:

```
WKLD(GROUP)
CYCLE(100)
DEVICE(COMM,NOUNITR,TAPE,DASD)
NOEXITS
OPTIONS
NOPAGESP
```

Because all options have not been filled in, input merge uses the program defaults (indicated earlier in this section under "Monitor I Session Options") to complete the list. Figure 3-3 shows the final option list, including the source for each option.

```
WKLD(GROUP) -- COMMAND          DEVICE(COMM) -- COMMAND
NOSYNC -- EXEC                  DEVICE(DASD) -- DEFAULT
NOTRACE -- MEMBER               DEVICE(TAPE) -- MEMBER
NOEXITS -- MEMBER               DEVICE(NOUNITR) -- MEMBER
NOENQ -- DEFAULT                CHAN -- DEFAULT
OPTIONS -- MEMBER               PAGING -- DEFAULT
STOP(8H) -- DEFAULT             CPU -- DEFAULT
CYCLE(100) -- COMMAND           MEMBER(02) -- EXEC
INTERVAL(30M) -- DEFAULT        MEMBER(10) -- COMMAND
RECORD -- DEFAULT               IOQ(NOCHRDR) -- DEFAULT
SYSOUT(A) -- DEFAULT            IOQ(NOCOMM) -- DEFAULT
NOREPORT -- DEFAULT             IOQ(DASD) -- DEFAULT
NOPAGESP -- MEMBER              IOQ(NOGRAPH) -- DEFAULT
DEVICE(NOCHRDR) -- DEFAULT      IOQ(NOTAPE) -- DEFAULT
DEVICE(NONMBR) -- DEFAULT       IOQ(NOUNITR) -- DEFAULT
DEVICE(NOGRAPH) -- DEFAULT      VSTOR(S) -- DEFAULT
DEVICE(NOSG) -- DEFAULT
```

Figure 3-3. Example of Input Merge for the Monitor I Session

When the Session Options Are Modified

When the MODIFY session command is used to modify the options established for an active session, RMF performs an input merge process. However, the process differs from the input merge used at the start of the session. Input merge as a result of a MODIFY session command starts with the list of options previously established and uses the input sources to **override** any previously established option.

The input sources and the order in which they are processed are:

1. The **options** field of the MODIFY session command. Any options specified here will override and replace any options in effect prior to the MODIFY session command.

2. The library data source. If the **options** field of the MODIFY command includes a MEMBER option, RMF will generate the member name and use any options specified in the member to **override** and replace any options in effect prior to the

processing of the library data source, including any options specified in the MODIFY command **options** field.

When more than one member is specified, RMF processes the members in left to right order. However, because the merge process that occurs in response to a MODIFY session command uses the input sources to override any previously-established options, options from the rightmost member will override any corresponding options from a previously-processed member. For example, if MEMBER(02,07) is specified on a MODIFY command, RMF generates the member names ERBRMF02 and ERBRMF07 and takes options from ERBRMF02 first. Thus, if ERBRMF02 specifies ENQ(DETAIL), the merge process places the option in the list of options currently established for the session. If, however, member ERBRMF07 specifies ENQ(SUMMARY), ENQ(SUMMARY) overrides the currently-established ENQ(DETAIL). ENQ(SUMMARY) is thus the option established for the session as a result of the MODIFY command.

When input merge is complete, RMF displays the options resulting from the merge at the operator console when an error is detected, when OPTIONS is in effect, or when conflicting options have been detected and changed. The actions RMF takes in response to errors and the response the operator can make are identical to those for input merge at the start of the session.

Example of Modifying Session Options

Assume that the options for a currently-active session include CHAN, NOCPU, and NOSTOP, and that you want to modify these options to NOCHAN, CPU, and STOP(40M).

If you issue the command:

```
MODIFY RMF.A,MODIFY ZZ,NOCHAN,CPU,STOP(40M)
```

the options will be modified as you want.

If, however, member ERBRMF10 includes the options:

```
NOCHAN
CPU
NOSTOP
```

and you issue the command:

```
MODIFY RMF.A,MODIFY ZZ,STOP(40M),MEMBER(10)
```

RMF will:

1. Merge the input option from the command and replace NOSTOP in the current option list with STOP(40M).
2. Merge the options from ERBRMF10 with the current options list, replacing CHAN with NOCHAN, NOCPU with CPU, and STOP(40M) with NOSTOP.

Thus, any options in a member will override both any current options **and** any options specified on the MODIFY session command.

 LC28-1556-2

Starting the Monitor I Session

The Monitor I session can be started concurrently with the initialization of RMF, as described in Chapter 2 under "RMF Initialization." Once RMF has been initialized, you can also start the Monitor I session by issuing a START session command.

The START session command is passed to RMF through the input field of the system MODIFY command. The required syntax is:

```
{MODIFY} {[RMF.]id,} {START} ZZ, [parm]
{  F   } {jobname, } {  S  }
```

id
The identifier assigned to RMF when it was initialized.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

ZZ
The identifier for the session to be started. For a Monitor I session, ZZ **must** be specified.

parm
The options for the session. Each option has the form

```
option[(value)]
```

Each option must be separated from the other options by a comma. If the MEMBER option is specified, any members identified in the **value** field are used during input merge to produce the list of options for the session. Options specified in this field of the START command have a higher priority than options in a library member.

When the START command is issued, RMF merges the input options. If the merge is free of error, RMF issues a message to indicate that the Monitor I session is active, and begins session processing.

There are three conditions that will cause the option list to be displayed at the operator console. They are:

- when OPTIONS is in effect
- when errors are detected during input merge
- when RMF has detected and changed conflicting options

At this point the operator can change the options, if necessary, before session processing begins.

Examples of Starting a Monitor I Session

1. To start a Monitor I session when all options are to be taken from other input sources, issue the command:

   ```
   MODIFY RMF.A,START ZZ
   ```

2. To start a Monitor I session that does not measure processor activity, produces reports at each interval, and samples data every 500 milliseconds, issue the command:

   ```
   MODIFY RMF.A,START ZZ,NOCPU,
         REPORT(REALTIME),CYCLE(500)
   ```

Modifying the Monitor I Session Options

The options in effect for an active Monitor I session can be modified by issuing a MODIFY session command specifying the options to be changed.

Issuing the MODIFY session command forces an immediate end of interval. After interval processing, RMF stops the session, modifies its options, and restarts the session. Although RMF automatically stops and restarts the session, you should be aware of these actions because no data is sampled during the brief period required to stop, modify, and restart the session.

A message is issued when the options specified on the MODIFY command have been changed and session processing has begun using the new options.

The MODIFY session command is passed to RMF through the input field of the system MODIFY command. The required syntax is:

```
{MODIFY} {[RMF.]id,} {MODIFY} ZZ[,options]
{  F   } {jobname,  } {  F   }
```

id
The identifier assigned to RMF when it was initialized.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

ZZ
The identifier for the session to be modified. For a Monitor I session, ZZ must be specified.

options
The session options to be changed. Each option has the form

```
option[(value)]
```

and must be separated by a comma. All options except SYSOUT can be specified; the SYSOUT option cannot be changed during a session. Note that, if you specify the MEMBER option, any options within the member will override any options specified in the command.

Example of Modifying Monitor I Session Options

To modify the options of an active Monitor I session to end the measurement of processor activity and magnetic tape device activity, issue the command:

```
MODIFY RMF.A,MODIFY ZZ,NOCPU,DEVICE(NOTAPE)
```

Note: At the start of the session, RMF uses the values in the INTERVAL and STOP options to calculate the number of measurement intervals for the session. If the STOP value is 30 minutes and the INTERVAL value is 30 minutes, there is one measurement interval during the session. When you change either of these values during the session, the duration of the session can be affected. Assume that a session has been started with the STOP and INTERVAL values noted above. After 15 minutes, you decide to change INTERVAL to 15 minutes, while STOP remains at 30 minutes. RMF re-calculates the number of intervals required by dividing the STOP value by the INTERVAL value. In this example, the number of intervals would be two. The duration of the session is therefore 45 minutes.

Ending the Monitor I Session

The Monitor I session can be ended in three ways:

- Issuing a system STOP command that stops all active non-TSO RMF sessions. The syntax of the command and an example of its use are given earlier in this book under "RMF Termination" on page 2-10.
- Expiration of the time value specified in the STOP option when the session was started. Use of the STOP option is described under "Monitor I Session Options" on page 3-1.
- Issuing a STOP session command to stop the Monitor I session. All other active sessions continue processing. This method is described here.

The STOP session command is passed to RMF through the input field of the system MODIFY command. The required syntax is:

```
{MODIFY} {[RMF.]id,} {STOP} ZZ
{  F   } {jobname, } { P  }
```

id
The identifier assigned to RMF when it was initialized.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

ZZ
The identifier for the session to be stopped. For a Monitor I session, ZZ **must** be specified.

Issuing this command forces an immediate end of interval; after interval processing is complete, RMF issues a message and ends the session.

Example of Ending the Monitor I Session

To stop the Monitor I session while allowing all other active RMF sessions to continue processing, issue the command:

```
MODIFY RMF.A,STOP ZZ
```

RMF Enqueue Contention Buffers

RMF initializes two 10,240-byte buffer areas from subpool 245 in order to collect global and local enqueue contention data. The module ERBMFIEQ contains the values that RMF uses to initialize the buffers. The offsets and the values are:

Offset	Name	Value	Description
100	GDATASZ	x'00002800'	global buffer size
104	LDATASZ	x'00002800'	local buffer size

If RMF collects more contention data than the buffer can hold, RMF issues the following message:

```
TABLE FULL-USE SPECIFIC NAME OR SHORT INTERVAL
```

If you receive this message, you should increase the size of the buffer; you can modify the size by means of the AMASPZAP program. An example follows:

```
//MODIFY      JOB
//STEP        EXEC        PGM=AMASPZAP
//SYSPRINT    DD          SYSOUT=A
//SYSLIB      DD          DSN=SYS1.LINKLIB,DISP=OLD
//SYSIN DD    *
    NAME      ERBMFIEQ    ERBMFIEQ
    VERIFY    0100        00002800
    VERIFY    0104        00002800
    REPL      0100        00002700
    REPL      0104        00002900
```

The statements in this example, reduce the global buffer size to 9,984 bytes and increase the local buffer size to 10,496 bytes. The minimum buffer size that RMF allows is 1,000 bytes (x'3E8'); that is, if you specify a value less than 1,000, RMF uses 1,000. There is no maximum buffer size. However if you specify a buffer size greater than 45,056 bytes (x'B000'), RMF fills the buffer but the data is truncated when passed to the SMF record. For more information on the use of AMASPZAP, see *Service Aids*.

Monitor I Session Interval Reports

This section describes the following Monitor I session interval reports:

- ASM/RSM/SRM trace activity
- Channel path activity
- CPU (processor) activity
- Enqueue activity
- I/O device activity
- I/O queuing activity
- Page/swap data set activity
- Paging activity
- Virtual storage activity
- Workload activity

At least one sample of each type of report is included in the following sections. Each entry in each report is described, including a field definition, its use, the range of values, and the field heading. All calculated numeric values in all reports are rounded to the nearest printable value, unless otherwise noted in the individual report descriptions. All data fields in the reports are obtained directly from the corresponding SMF record image unless otherwise indicated. Those data fields that are not obtained directly are calculated from fields in the SMF record image.

Monitor I session interval reports can be obtained during RMF processing or at the end of RMF processing, or they can be generated at a later time by the post processor.

Heading Information

Figure 3-4 is an example of the heading on all Monitor I session interval reports.

```
                              R E P O R T    T I T L E

                                                                          PAGE nnnn
      OS/VSn             SYSTEM ID cccc         DATE mm/dd/yy         INTERVAL mm.ss.ttt
      RELEASE            RPT VERSION 3.5.1      TIME hh.mm.ss         CYCLE s.ttt SECONDS
```

Figure 3-4. Example of a Heading on a Monitor I Session Report

All Monitor I session reports contain the following heading information:

Figure 3-5. Information in the Heading of Monitor I session reports

Heading	Definition
Report title	The identification of the type of measurement data.
OS/VSn	The type of operating system, where n is the system indicator obtained from the system indicator field in the SMF record.
SYSTEM ID cccc	The four-character SMF identifier associated with this system at system generation, where cccc is the identifier.
RELEASE	The release number and level of the operating system.
RPT VERSION v.l.m	The three-digit RMF report level, where v is the device version number, l is the level number, and m is the modification level number.
DATE mm/dd/yy	The date of the beginning of the interval for which the measurement was gathered, where mm is the month, dd is the day, and yy is the year.
INTERVAL mm.ss.ttt	The length of the interval at which unique sets of measurements are gathered for input to the report generators,where mm is the minutes, ss is seconds, and ttt is thousandths of seconds.
TIME hh.mm.ss	The time the interval began, where hh is hours, mm is the minutes, and ss is seconds.
CYCLE s.ttt	The length of the cycle at which data is sampled, where s is seconds and ttt is thousandths of seconds. This field appears in the heading for each report that includes sampled data.
PAGE nnnn	The page number of the report (generated by the report program), where nnnn is the page number.

ASM/RSM/SRM Trace Activity Report

The ASM/RSM/SRM trace activity report provides information about certain fields of the RSM control and enumeration area (RCE), the auxiliary storage manager vector table (ASMVT), the system resource management data areas, and the SRM domain tables and logical path blocks (LPBs). The fields that RMF can trace are listed earlier in this chapter under "Monitor I Session Options." RMF trace treats values collected as unsigned binary integers. In addition to these fields, you can add other fields to the report, as described in "Monitor I Session User Reports" in Chapter 7.

The trace report shows snapshots of each of the specified field values along with timing information. One important use of the trace report is to monitor the SRM multiprogramming level (MPL) adjustment. You can monitor the contention SRM detects and see the actions SRM takes as a result of contention by tracing the following fields in the SRM data area and domain table: RCVUICA, RCVCPUA, RCVASMQA, RCVPTR, RCVTOTDF, DMDTMPLT, DMDTRUA, and DMDTCMPL.

On a 3090 processor, you can monitor the use of extended storage by tracing the RCE fields RCEESINU, RCEESWRT,RCENWSF and RCEWSDNE, and the SRM fields MCVMGAGE and MCVSTCRI.

Tracing the domain fields enables you to trace the multiprogramming target level over the interval and compare it with the current multiprogramming level that SRM is aiming at for each domain. Thus, by tracing the domain fields, you can determine whether the domain setting in the IPS is appropriate. Domain tracing ends on a SET IPS operator command and will not restart unless the session options are modified to re-specify domain tracing or the Monitor I session is stopped and restarted. When this occurs, **ENDED** appears on the report for the domain.

Tracing logical path blocks (LPBs) enables you to trace each LPB being monitored by the SRM. Each of these traces has the heading "LPBUTIL" followed by a list of the channel paths included in that logical path block.

The options that can be selected for a given report on a given traceable item include:

Figure 3-6. Traceable Options

Option	Meaning
MIN	The minimum value of the field over the period of sampling.
MAX	The maximum value of the field over the period of sampling.
AVG	The average value of the field.
END	The value of the field when the last sample was taken.
STDDEV	The standard deviation of the values sampled.
ALL	The default, meaning all of the above options.

See TRACE in "Monitor I Session Options" for how to specify these options.

The number of lines in the report is based on the cycle and interval values specified when the session is started. For example, if CYCLE(250) and INTERVAL(60M) are specified, approximately 14,400 samples are taken. Based on a constant 60 samples/set, and 1 line/set, the report contains 240 lines of data for each field.

The number of samples per set is determined by a constant located in the first halfword of CSECT ERBMFTTB in load module ERBMFMFC (in SYS1.LINKLIB). Any value in the range 1 to 32,767 is valid and can be changed by the system programmer. If changed to zero, RMF overrides it with 1. If changed to a value less than 0 (a negative number), RMF defaults to 32,767 (X'7FFF'). If the values specified for cycle, interval, and samples per set would result in the number of sets exceeding this limit, RMF suspends trace sampling for the remainder of the interval.

Figure 3-7 shows a sample printed report for trace activity. Each field is defined in Figure 3-8 on page 3-29

```
                                              T R A C E   A C T I V I T Y
                                                                                                                        PAGE 1

          OS/VS2                        SYSTEM ID AQFT             DATE 03/04/89               INTERVAL 29.59.997
          SP2.2.3                       RPT VERSION 3.5.1          TIME 09.00.00               CYCLE 1.000 SECONDS

   SECONDS/SET=  60.00  CYCLES/SAMPLE=  1      NUMBER OF SAMPLES=  1800   SAMPLES/SET=    60   NUMBER OF SETS=  30

   TIME    * SMCABFLS * CCVUTILP * RCVAFQA  * RCVMSPP  * RCVPAGRT * RCVDPR   * RCVPTR   * MCVSBLTF * RCVMFXA  * RCVFXIOP *
 MM.SS.TT  *  AVERAGE *  AVERAGE *  AVERAGE *      END *      END *      END *      END *  AVERAGE *  AVERAGE *  AVERAGE *

 01.00.00  *     0.00 *    38.68 *  1738.31 *       13 *        3 *        0 *        0 *    22.00 *    22.08 *    16.00 *
 02.00.00  *     0.00 *    41.90 *   949.78 *        0 *        0 *        0 *        5 *    22.00 *    22.43 *    16.00 *
 03.00.00  *     0.00 *    63.03 *  1110.23 *       35 *        2 *        0 *        1 *    22.00 *    21.81 *    16.00 *
 04.00.00  *     0.00 *    55.05 *  1629.30 *       10 *        0 *        0 *        0 *    22.00 *    22.88 *    16.00 *
 05.00.00  *     0.00 *    41.46 *   788.26 *        0 *        0 *        0 *        0 *    22.00 *    23.00 *    16.76 *
 06.00.00  *     0.00 *    46.65 *   822.53 *       59 *       15 *        0 *        0 *    22.13 *    22.91 *    16.65 *
 07.00.00  *     0.00 *    45.25 *   632.25 *       15 *        1 *        0 *        0 *    22.05 *    22.91 *    16.23 *
 08.00.00  *     0.00 *    49.65 *   987.48 *       17 *        4 *        0 *        0 *    22.00 *    22.58 *    16.40 *
 09.00.00  *     0.00 *    43.68 *  1584.30 *       34 *        3 *        0 *        1 *    22.00 *    23.00 *    16.91 *
 10.00.00  *     0.00 *    33.48 *   671.20 *        0 *        0 *        0 *        0 *    22.00 *    22.91 *    17.00 *
 11.00.00  *     0.00 *    28.78 *   610.65 *       15 *        5 *        0 *        0 *    22.00 *    23.00 *    17.00 *
 12.00.00  *     0.00 *    39.78 *   659.65 *       34 *        1 *        0 *        0 *    22.00 *    22.65 *    16.56 *
 13.00.00  *     0.00 *    32.21 *  1522.53 *        0 *        0 *        0 *        0 *    22.03 *    23.00 *    17.00 *
 14.00.00  *     0.00 *    47.83 *   605.83 *       72 *       58 *        0 *        0 *    23.00 *    23.16 *    17.00 *
 15.00.00  *     0.00 *    37.78 *   866.43 *        0 *        0 *        0 *        0 *    23.00 *    23.95 *    17.00 *
 16.00.00  *     0.00 *    36.58 *   969.85 *       39 *       27 *        0 *        3 *    22.98 *    23.60 *    16.68 *
 17.00.00  *     0.00 *    54.90 *  4010.61 *       52 *       21 *        0 *        0 *    22.00 *    22.31 *    16.13 *
 18.00.00  *     0.00 *    43.20 *  1199.55 *       52 *      126 *        0 *        0 *    22.65 *    23.00 *    17.00 *
 19.00.00  *     0.00 *    56.51 *   813.15 *       43 *      134 *        0 *        0 *    23.00 *    23.16 *    17.00 *
 20.00.00  *     0.00 *    65.88 *  1079.26 *       24 *       30 *        0 *        0 *    23.10 *    23.65 *    17.00 *
 21.00.00  *     0.00 *    54.16 *  1227.21 *       38 *       14 *        0 *        6 *    23.00 *    23.91 *    17.00 *
 22.00.00  *     0.00 *    59.98 *   824.30 *       42 *      129 *        0 *        5 *    23.00 *    24.00 *    17.00 *
 23.00.00  *     0.00 *    51.85 *  1108.93 *       46 *       78 *        0 *        0 *    23.00 *    23.83 *    17.00 *
 24.00.00  *     0.00 *    48.60 *  1029.20 *       44 *       43 *        0 *        0 *    23.00 *    23.50 *    17.00 *
 25.00.00  *     0.00 *    47.65 *  1234.71 *       31 *        1 *        0 *        0 *    23.00 *    24.00 *    17.00 *
 26.00.00  *     0.00 *    51.91 *  1135.13 *        0 *        0 *        0 *        0 *    23.38 *    24.16 *    17.00 *
 27.00.00  *     0.00 *    49.01 *  1349.71 *        0 *        1 *        0 *        0 *    24.00 *    24.30 *    17.00 *
 28.00.00  *     0.00 *    52.01 *   742.91 *       50 *      201 *        0 *        0 *    24.00 *    24.78 *    17.00 *
 29.00.00  *     0.00 *    55.86 *  1328.58 *       71 *        7 *        0 *        0 *    24.16 *    24.43 *    17.00 *
 30.00.00  *     0.00 *    60.25 *  2006.26 *       36 *        4 *        0 *        0 *    24.46 *    25.00 *    17.23 *
   MAXIMUM*     0.00 *    65.88 *  4010.61 *       72 *      201 *        0 *        6 *    24.46 *    25.00 *    17.23 *
   MINIMUM*     0.00 *    28.78 *   605.83 *        0 *        0 *        0 *        0 *    22.00 *    21.81 *    16.00 *
   AVERAGE*     0.00 *    47.78 *  1174.60 *    29.06 *    30.26 *     0.00 *     0.70 *    22.69 *    23.33 *    16.75 *
```

Figure 3-7. ASM/RSM/SRM Trace Activity Report

Figure 3-8. Fields in ASM/RSM/SRM Trace Activity Report

Heading	Definition
SECONDS/SET	Amount of elapsed time covered by one line of output
CYCLES/SAMPLE	Number of cycles in a sample
NUMBER OF SAMPLES	Total number of samples taken over the interval
SAMPLES/SET	Number of samples taken for each line of output except the last line; it can contain fewer samples
NUMBER OF SETS	Number of output lines
TIME / MM.SS.TT	The approximate calculated time when the sampling for that line of data ended (minutes, seconds and thousandths of a second).
MINIMUM	The smallest value sampled for the period covered by that output line.
AVERAGE	The average of the values collected for the period covered by that output line.
MAXIMUM	The largest value sampled for the period covered by that output line.
END	The last value sampled for the period covered by that output line.
STD. DEV.	The standard deviation of the values collected for the period covered by that line of output.

Standard Deviation Output

The trace activity report gives an account of standard deviation as an E-format (exponential) number. This number expresses the standard deviation for a set (sub-interval) as a number between 0.000 and 9.999, times a power of ten.

Example 1

```
1.123E +01 is the same as 1.123 X 10.1. or 11.23
```

Note: The signed number following the E represents a power of 10 and indicates that the decimal point should be shifted either to the left or right. In this example, the decimal point is moved one place to the right.

Example 2

```
2.903E -03 or .002903
```

Note: In this example -03 is the same as 10-.3.. The decimal point is moved three places to the left.

Example 3

```
6.031E +00 or 6.031
```

Note: In this example +00 is the same as 10.0., which is equal to 1. Therefore, the decimal point should not be shifted at all.

Domain Table Tracing

Tracing the domain table fields causes the domain numbers to be appended to an abbreviation of the field name on the report:

Field Specified	Field Reported
DMDTINCU	DMINUxxx
DMDTOUTU	DMOUTxxx
DMDTMPLT	DMTARxxx
DMDTRUA	DMRUAxxx
DMDTCMPL	DMMPLxxx
DMDTTWSR	DMTWSxxx
DMDTCIDx	DMCDXxxx
DMDTTWET	DMTWExxx

Each field name specified is traced and reported for all domains that were active when RMF was initialized. For instance, if domains 0, 1, 46, and 245 are active when RMF starts, and DMDTINCU is the field requested, the following names are presented for the options requested:

DMINU000
DMINU001
DMINU046
DMINU245

Channel Path Activity Report

The channel path activity report provides information about channel path use. The channel path is the physical interface that connects control units and I/O devices to the I/O subsystem. Channel path activity information, when used with I/O device activity and I/O queuing activity information, can be used to identify performance bottlenecks associated with channel paths. For example, if a channel path to a device shows excessive use, you could take steps to remedy the situation, such as defining additional paths to the device or introducing a different job mix to produce better performance.

The report identifies each channel path by identifier and channel path type and indicates the percentage of time when the channel path was busy. Because the maximum number of channel paths that RMF can report on a single page is 96, channel path activity is usually a single page report. On a 3090 processor with more than 96 channel paths, RMF generates a two-page report.

The report includes data for each valid (enabled) channel path. Data, however, does not appear for any channel path that was offline at the end of the interval or that was brought online during the interval. Instead, one of the following messages appears in the data field:

- NOW ONLINE -- Brought online during the interval and still online at the end of the interval
- NOW OFFLINE -- Taken offline during the interval and still offline at the end of the interval
- OFFLINE -- Offline for the entire interval

Duration Report: When you request a duration report for channel path activity, the identifier of each channel that was varied online or offline during the reporting period is followed by a single asterisk (*).

Figure 3-9 shows a sample printed report for channel path activity. Each field is defined in Figure 3-10 on page 3-33.

C H A N N E L   P A T H   A C T I V I T Y

PAGE 1

OS/VS2	SYSTEM ID AQFT	DATE 03/04/89	INTERVAL 29.59.997
SP2.2.3	RPT VERSION 3.5.1	TIME 09.00.00	CYCLE 1.000 SECONDS

CHANNEL PATH ID	TYPE	PERCENT CH PATH BUSY	CHANNEL PATH ID	TYPE	PERCENT CH PATH BUSY	CHANNEL PATH ID	TYPE	PERCENT CH PATH BUSY	CHANNEL PATH ID	TYPE	PERCENT CH PATH BUSY
00	BY	0.00	08	BL	0.00	10	BL	0.00	18	BL	1.16
01	BL	0.00	09	BL	15.02	11	BL	15.10	19	BL	4.31
02	BL	0.56	0A	BL	0.09	12	BL	0.00	1A	BL	11.25
03	BL	9.86	0B	BL	0.00	13	BL	0.00	1B	BL	6.40
04	BL	15.83	0C	BL	0.00	14	BL	4.29	1C	BL	58.95
05	BL	0.12	0D	BL	0.97	15	BL	0.00	1D	BL	3.54
06	BL	0.00	0E	BL	2.07	16	BL	0.53	1E	BL	9.11
07	BL	0.00	0F	BL	0.21	17	BL	10.77	1F	BL	0.00
20	BL	0.00	28	BL	4.58	40	BL	0.00	48	BL	0.00
21	BL	0.04	29	BL	0.00	41	BL	0.00	49	BL	12.70
22	BL	0.00	2A	BL	0.00	42	BL	0.54	4A	BL	0.01
23	BL	0.00	2B	BL	0.05	43	BL	5.41	4B	BL	0.00
24	BL	7.23	2C	BL	1.53	44	BL	0.00	4C	BL	0.00
25	BL	0.00	2D	BL	0.03	45	BL	0.00	4D	BL	0.00
26	BL	0.00	2E	BL	7.99	46	BL	0.00	4E	BL	2.46
27	BL	0.00	2F	BL	10.16	47	BL	0.00	4F	BL	0.00
50	BL	0.00	58	BL	2.12	60	BL	0.00	68	BL	4.66
51	BL	14.54	59	BL	8.51	61	BL	0.05	69	BL	0.00
52	BL	0.00	5A	BL	7.97	62	BL	0.00	6A	BL	14.37
53	BL	0.00	5B	BL	9.10	63	BL	0.00	6B	BL	0.04
54	BL	1.91	5C	BL	7.35	64	BL	1.04	6C	BL	2.43
55	BL	0.00	5D	BL	4.06	65	BL	8.28	6D	BL	0.49
56	BL	0.55	5E	BL	5.64	66	BL	0.00	6E	BL	12.43
57	BL	5.42	5F	BL	0.00	67	BL	0.00	6F	BL	0.00

Figure 3-9. Channel Path Activity Interval Report

Figure 3-10. Fields in Channel Path Activity Report

Heading	Definition
CHANNEL PATH ID	The hexadecimal number of the channel path identifier (CHPID).
TYPE	Either BY for byte multiplexor or BL for block multiplexor. **For a 308x, 908x, or 3090 processor**, RMF uses the I/O configuration data set (IOCDS) to find the channel path type. If RMF cannot read the data set, or if the data set has been updated so that the data might not apply to the present configuration, RMF issues the status message IOCDS INFORMATION UNAVAILABLE TO RMF. Two different response codes are issued: • A response code of 2024 indicates that the IOCDS has been updated so that the data might not apply to the present configuration. • A response code of 43F0 indicates that RMF cannot read the IOCDS because it appears in another partition of a multi-processing system (for example, when RMF is running on a 3084 processor, and the operator partitions the system in such a way that RMF cannot read the IOCDS). In either case, RMF cannot obtain the channel path type. RMF continues to measure channel path activity, but the requested SMF record does not contain the channel path type, and the field remains blank in the channel path activity report. **For a 4381 processor complex**, RMF uses the I/O configuration data (IOCD) in the hardware system area (HSA) to retrieve the channel path type. If RMF cannot retrieve the configuration data after successive attempts, RMF does not report the channel path type; the field is blank.
PERCENT CH PATH BUSY	The percentage of time during the interval when the channel path was busy. To calculate the value, RMF uses the sampled data the SRM collects in the channel path measurement table (CPMT). The calculation used is: `%CH PATH BUSY = 100 x (SB/N)` The percentage equals 100 times the number of SRM observations of the channel path busy (SB), divided by the number of SRM samples taken (N).

Central Processor Unit (CPU) Activity Report

The central processor unit (CPU) activity report provides processor information followed by a section headed SYSTEM ADDRESS SPACE ANALYSIS that provides overall address space information. In a PR/SM environment in LPAR mode, RMF also provides a partition data report.

The processor information section of the report is headed by the CPU MODEL field, which identifies the processor by model number and version of that model.

For each processor, the report provides the following information:

- Identification of the processor number and serial number
- Vector Facility (VF) indicator and vector affinity percentage
- Processor busy time expressed as a percentage of the interval
- Processor wait time expressed as a percentage of the interval
- Interrupt rate (the rate of I/O interrupts per second during the measurement interval, including interrupts handled by the second level interrupt handler (SLIH) and those handled by the Test Pending Interrupt (TPI) instruction)
- Percentage of I/O interrupts that were handled by the TPI instruction

Following the data lines for each processor is a single line that reports the VF affinity percentage, the average busy time percentage, the average wait time percentage, the total I/O interrupt rate, and the average TPI interrupt percentage, over all processors.

Data lines are included for each processor that was online at the end of at least one interval. However, data is not formatted for a processor that was offline during the entire interval or that came online during the interval. Processors for which data is not printed have one of the following messages in the first data field:

- NOW ONLINE - Brought online during the interval and still online at the end of the interval
- NOW OFFLINE - Taken offline during the interval and still offline at the end of the interval
- OFFLINE - Offline for the entire interval

Duration Report: When you request a duration report for processor activity, the identifier of each processor that moved online or offline during the reporting period is followed by a single asterisk (*).

The heading for the address space analysis section includes the field SAMPLES, which describes the total number of samples taken during the RMF reporting interval.

The data analyzes the following types of address spaces:

- In storage and ready to execute
- In storage
- Out of storage and ready to execute
- Out of storage and not ready to execute
- Logically out of storage and ready to execute
- Logically out of storage and not ready to execute.

Data is also presented on the number of address spaces used by batch users, by started tasks and mount tasks, and by TSO users. Examining the data can indicate when a backlog of address spaces are waiting to use the processor.

Figure 3-11 shows a sample printed report for CPU activity. Each data field is defined in Figure 3-12 on page 3-36.

C P U   A C T I V I T Y

PAGE 1

OS/VS2	SYSTEM ID AQFT	DATE 03/22/89	INTERVAL 29.59.957
SP2.2.3	RPT VERSION 3.5.1	TIME 10.00.00	CYCLE 1.000 SECONDS

CPU MODEL   3090   VERSION 61

CPU NUMBER	VF ONLINE	VF AFFINITY PERCENTAGE	BUSY TIME PERCENTAGE	WAIT TIME PERCENTAGE	CPU SERIAL NUMBER	I/O TOTAL INTERRUPT RATE	% I/O INTERRUPTS HANDLED VIA TPI
0	---	****	96.19	----	010028	29.30	9.86
1	---	****	96.11	----	110028	27.24	11.14
2	---	****	96.00	----	210028	1699	17.03
TOTAL/AVERAGE		****	96.10	0.00		1756	16.81

SYSTEM ADDRESS SPACE ANALYSIS   SAMPLES = 1,800 .

TYPE	NUMBER OF ASIDS MIN	MAX	AVG	DISTRIBUTION OF QUEUE LENGTHS (%) 0	1	2	3	4	5	6	7-8	9-10	11-12	13-14	14+
IN READY	0	33	10.8	0.2	0.6	2.0	3.7	4.5	5.5	6.2	15.0	14.0	12.0	10.8	25.1

TYPE	MIN	MAX	AVG	0	1-2	3-4	5-6	7-8	9-10	11-15	16-20	21-25	26-30	31-35	35+
IN	56	75	64.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0
OUT READY	0	4	0.1	92.1	7.1	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
OUT WAIT	102	137	121.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0
LOGICAL OUT RDY	0	9	0.1	90.1	8.6	0.6	0.3	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
LOGICAL OUT WAIT	19	48	33.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	7.7	21.0	40.0	30.9
BATCH	10	24	18.3	0.0	0.0	0.0	0.0	0.0	0.8	20.8	44.3	33.8	0.0	0.0	0.0
STC	61	79	67.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0
TSO	130	137	133.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0

Figure 3-11. CPU Activity Report

Figure 3-12 (Page 1 of 2). Fields in the CPU activity Report

Heading	**Definition**
CPU MODEL	The model number of the CPU.
VERSION	The CPU version number (hexadecimal).
CPU NUMBER	A number that indicates the internal address of the processor for which values are being reported.
VF ONLINE	The Vector Facility (VF) indicator for the processor. If the Vector Facility is installed on this processor and online during the entire interval, the column contains "YES"; otherwise, the column contains "--."
VF AFFINITY PERCENTAGE	The percentage of vector affinity time for this processor. For more information on the Vector Facility affinity time, see *MVS/XA System Programming Library System Management Facilities (SMF)*. When the Vector Facility is indicated for the processor (VF ONLINE), the calculation is: `VAP = 100 * (VFS/S)` The vector affinity percentage (VAP) equals 100 times the number of samples for which the Vector Facility bit was found set (VFS) divided by the total number of samples in the interval (S). Eleven asterisks (***********) appear in this column when the Vector Facility is not indicated for the CPU, or when the Vector Facility is installed and online, but RMF cannot calculate the affinity percentage because no samples are available for the processor during the report interval.

BUSY TIME PERCENTAGE	The percentage of the interval time that the processor was busy.
	The calculation used is:
	`BTP = (Interval Time - Wait Time) / Interval Time X 100`
	Wait Time is explained in the following field.
	In a PR/SM environment in LPAR mode, the BTP calculation is:
	• For a dedicated partition:
	Use the above calculation.
	• For a non-dedicated partition with Wait Completion NO:
	`BTP = Processor Dispatch Time / Interval Time X 100`
	The processor dispatch time is the elapsed time that PR/SM dispatched this processor during the interval. Part of this dispatch time may be MVS system wait time.
	• For a non-dedicated partition with Wait Completion YES:
	`BTP = (Processor Dispatch Time - Wait time) / Interval Time X 100`

Figure 3-12 (Page 2 of 2). Fields in the CPU activity Report

Heading	Definition
WAIT TIME PERCENTAGE	The percentage of the reporting interval time during which the processor was not required by any user currently in the system or by the operating system itself. The Wait Time Percentage (WTP) calculation is: `WTP = (Wait Time / Interval Time) X 100` **Note:** A WTP value greater than 100 is probably caused by the operator having entered a SET CLOCK command to reset the TOD clock after the RMF session was started. If the wait time percentage field exceeds 100 percent, the busy time percentage field is zero. **In a PR/SM environment in LPAR mode,** the WTP calculation is: • For a non-dedicated partition with Wait Completion NO: An accurate wait time is not available. Dashes (----) will appear in this field. • In all other cases: `WTP = (Wait Time / Interval Time) X 100` Following the data lines for each processor, the report lists the average wait time for all processors. TOTAL/AVERAGE in the CPU number column identifies the TOTAL/AVERAGE value as the average of all wait time percentages reported.
CPU SERIAL NUMBER	The six-digit number that matches the physical serial number stamped on the processor that, along with the model number, uniquely identifies the processor. In a PR/SM environment in LPAR mode, the second character of the CPU serial number is the PR/SM logical partition definition (LDEF) panel partition id.
I/O TOTAL INTERRUPT RATE	The total rate per second at which this processor handled I/O interruptions. The rate reported reflects the processing for the entire interval, which might include periods of time when the SRM enabled or disabled this processor for I/O interruptions. The rate includes interruptions that the second level interrupt handler (SLIH) handled as well as those the Test Pending Interrupt (TPI) instruction handled. The calculation used is: `R = (SLIH+TPI)/T` The total I/O interruption rate (R), equals the number of interruptions the SLIH handled plus the number of interruptions the TPI instruction handled, divided by the length of the reporting interval (T). Following the data lines for each processor, the average I/O interruption rate for all processors included is reported. This line is identified by the label TOTAL/AVERAGE in the CPU number column. The TOTAL/AVERAGE is the average of all interruption percentages reported.
% I/O INTERRUPTS HANDLED VIA TPI	The percentage of the total interrupts for this processor during the RMF interval that the TPI instruction handled. The calculation used is: `P = (100 x TPI)/(SLIH+TPI)` The percentage (P), equals 100 times the number of interruptions handled by the TPI instruction divided by the total number of interruptions. Following the data lines for each processor, the average number of interruptions for all processors included is reported. This line is identified by the label TOTAL/AVERAGE in the CPU number column. The TOTAL/AVERAGE value is the weighted average (weighted by I/O interrupt rate) of all interruption percentages reported.

Figure 3-13. Fields in the Address Space Analysis Section

Heading	Definition
TYPE	The following types of address spaces are analyzed: IN READY — Address spaces that are in real storage and ready to execute. IN — Address spaces that are in real storage (corresponds to SRM in queue). This count includes the IN READY count. OUT READY — Address spaces on the SRM out queue that are physically swapped out of real storage and ready to execute. **Note:** Some address spaces on the SRM out queue might represent those TSO users that the SRM intentionally delayed to meet an installation's response time objective. Because these address spaces do not represent a potential performance problem, they are not included in the value reported for OUT READY. OUT WAIT — Address spaces on the SRM wait queue that are physically swapped out of real storage and not ready to execute. LOGICAL OUT RDY — Address spaces on the SRM out queue that are physically in real storage but logically swapped out of real storage and ready to execute. LOGICAL OUT WAIT — Address spaces on the SRM wait queue that are physically in real storage but logically swapped out of real storage and not ready to execute. BATCH — Address spaces used for batch jobs. STC — Address spaces used for started tasks or mount tasks. TSO — Address spaces used for TSO users.
NUMBER OF ASIDS	Number indicating the quantity of ASIDs of each type detected during the reporting interval, reported as: MIN — The minimum number of address spaces for each type MAX — The maximum number of address spaces for each type AVG — The average number of address spaces for each type
DISTRIBUTION OF QUEUE LENGTH (%)	The percentage of samples during the reporting interval when the number of address spaces of a particular type fell within a certain range of values. The range of values is indicated in the subheadings.

The **partition data report** section of the CPU activity report provides data on all configured partitions. This section only appears when RMF is running in a PR/SM environment in LPAR mode.

Each partition created in LPAR mode must have RMF installed in order to measure the performance of the MVS system running in that partition. If MVS is running as a second level operating system in a partition, for example under VM, RMF will report only on the MVS system.

The partition data report section of the CPU activity report is divided into three parts:

- Header information (first four lines)
- Logical partition data
- Logical partition processor data.

The header information gives an overview of the LPAR mode characteristics. These four lines appear on all partition data reports.

The logical partition data section provides the following information for each partition:

- Name
- Status
- Weighting share of resources
- Wait completion option.

The logical partition processor data part of the partition data report, provides the following information about the partition's processors:

- Number of logical processors assigned to this partition
- The partition's total processor dispatch time
- The partition's average processor dispatch time
- The partition's average logical processor utilization
- The partition's physical processor utilization.

Note: If an asterisk (*) appears next to the following fields, care should be taken when interpreting data in related fields:

- NUMBER OF PHYSICAL PROCESSORS
- DISPATCH INTERVAL
- STATUS
- WEIGHTS
- WAIT COMPLETION
- NUMBER OF LOG PRCRS.

An * indicates a change has been made to the value during the measurement interval.

Figure 3-14 on page 3-40 shows an example of the partition data report section of the CPU activity report. This section appears only when running in a PR/SM environment in LPAR mode.

```
                              P A R T I T I O N   D A T A   R E P O R T
                                                                                 PAGE    2
        OS/VS2                     SYSTEM ID MVS2          DATE 03/14/89        INTERVAL 01.00.000
        SP3.1.0                    RPT VERSION 3.5.1       TIME 20.57.13        CYCLE 1.000 SECONDS

MVS PARTITION NAME                    PARTN 1
NUMBER OF CONFIGURED PARTITIONS          4
NUMBER OF PHYSICAL PROCESSORS            6
DISPATCH INTERVAL                     DYNAMIC

-------LOGICAL PARTITION DATA--------      ---------LOGICAL PARTITION PROCESSOR DATA---------
                            WAIT        NUMBER OF   ---DISPATCH TIME DATA---- AVRG LOG   PHYSICAL
  NAME   STATUS WEIGHTS COMPLETION      LOG PRCRS     TOTAL        AVERAGE    PRCR UTIL  PRCRS UTIL
 PARTN 1   A      200        NO             1      00.00.41.976 00.00.41.976   69.96      11.66
 PARTN 2   A      300        NO             2      00.01.03.000 00.00.31.500   52.50      17.50
 PARTN 3   D
 PARTN 4   A      500        NO             4      00.01.45.192 00.00.26.298   43.83      29.22
                                                   ------------                           ------
  TOTAL                                            00.03.30.168                           58.38
```

Figure 3-14. Partition Data Report Section of CPU Activity Report

Figure 3-15 (Page 1 of 2). Fields in the Partition Data Report

Heading	Definition
MVS PARTITION NAME	The name of the partition under which RMF generated this partition data report.
NUMBER OF CONFIGURED PARTITIONS	Total number of activated and deactivated configured partitions.
NUMBER OF PHYSICAL PROCESSORS	Number of physical processors in PR/SM system. An asterisk (*) next to this value indicates a change to this count during the measurement interval. The value at the end of the measurement interval is reported.
DISPATCH INTERVAL	The maximum number of milliseconds a processor is used when dispatched. The value applies only to partitions not running on dedicated processors. If an asterisk (*) is next to this value indicates a change in the dispatch interval. If DYNAMIC appears in this field, you did not specify a value to the set running time field in the PR/SM logical partition control (LPCTL) panel. PR/SM dynamically allocates how long a processor can be assigned to a partition.
NAME	One to eight character name that identifies a partition.
STATUS	Indicates current status of the partition. Either: • A - Activated • D - Deactivated. If a partition is deactivated, the rest of the report line is blank. An asterisk (*) next to this value indicates the partition status has changed during the measurement interval. The status at the end of the interval appears. If an asterisk appears for this field, the following three fields will also show asterisks.
WEIGHTS	The partition's current weighting of the shared processor resources. If the partition is dedicated, a share value is not assigned and DED will appear in this field. An asterisk (*) next to this value indicates the weight value has been changed during the measurement interval.
WAIT COMPLETION	Indicates whether the wait completion option is set to YES or NO. YES implies that the processors assigned to this partition will remain dispatched to the partition until the Dispatch Interval has expired. NO implies that the processors assigned to this partition become available to other partitions when the work for this partition is completed. The Dispatch Interval may or may not be expired. This field is blank for a dedicated partition. A processor will not remain dispatched to a partition for a period longer than the current dispatch interval, whether it is a fixed value or dynamically determined. An asterisk (*) next to this value indicates a change in this option during the measurement interval.
NUMBER OF LOG PRCRS	The number of logical processors assigned to this partition. An asterisk (*) next to this value indicates a change sometime during the measurement interval. You can change the value only by deactivating then reactivating the partition.

Figure 3-15 (Page 2 of 2). Fields in the Partition Data Report

Heading	Definition
TOTAL DISPATCH TIME	The sum of all processor dispatch times, for this partition, that occurred during the measurement interval; expressed in the form HH.MM.SS.TTT. Processor dispatch time is the time that a processor is available to this partition.

AVERAGE DISPATCH TIME	The average dispatch time per processor, expressed in the form HH.MM.SS.TTT. The calculation is:
	`           Total Dispatch Time`
	`AVERAGE = ---------------------`
	`           Number of Log Prcrs`
	Total Dispatch Time is defined in the field above.

AVRG LOG PRCR UTIL	The average logical processor utilization for this partition. The calculation is:
	`AVRG LOG        Average Dispatch Time`
	`PRCR UTIL = -------------------------  X 100`
	`                   Interval Time`
	This value includes any wait time that may have occurred in the system running in the partition. See the first section of the CPU activity report for a breakdown of the system wait time.

PHYSICAL PRCRS UTIL	This partition's utilization of the total PR/SM physical processor resource. The calculation is:
	`PHYSICAL     Total Dispatch Time For This Partition`
	`PRCRS UTIL = ------------------------------------------ X 100`
	`            NUMBER OF PHYSICAL PROCESSORS X Interval Time`
	This value is reported for partitions with shared or dedicated physical processors.

TOTAL	The sum of the total dispatch times for each processor represents the total amount of time the processor was available. The sum of the PHYSICAL PRCRS UTIL field represents the total utilization of the PR/SM physical processor resource by all partitions.

Enqueue Activity Report

The enqueue activity report provides information about resources that periodically build up queues of one or more requestors waiting to use the resource. Because the amount of time that a requestor must spend waiting for a resource can seriously affect system throughput, the information in this report can be very helpful in locating resources that consistently cause bottlenecks. Contention is reported for those resources to which access is controlled by jobs that issue ENQ and DEQ macro instructions.

Once you have defined a critical resource, such as a serially-reusable resource that can be requested on either an exclusive or shared basis, your installation can improve the situation in a variety of ways. You could change the hardware configuration to release device bottlenecks, change data set placement, or reschedule jobs to improve throughput, or re-specify the installation tuning parameter ERV (enqueue residence value) to give more processor time to the holder of the resource.

The report can be requested on two levels: summary and detail. The summary report includes all resources for which contention has occurred during the reporting interval, a description of the contention time for each resource, a queue length distribution and average queue length for each resource, information on the type of requests made (either exclusive or shared), and the total number of enqueue contention events that occurred. An enqueue contention event is defined as the period from the time when the resource first has contention until the resource no longer has contention. A resource for which contention is still occurring at the end of the interval will be indicated by an asterisk following the TOT field, which is under the CONTENTION TIME field.

The detail report includes additional information: (1) the total number of jobs that own the resource and the names of one or two jobs that own the resource and (2) the total number of jobs that are waiting for the resource and the names of one or two jobs that are waiting for the resource. RMF selects the job names shown in the detail report during the period of maximum contention in the interval by (1) determining the longest contention event in the interval and (2) for that event, reporting the owners and waiters at the point when the event queue is the longest. Note that when there are several occurrences of the same length queue, the latest queue is reported. Each job name is followed by (1) the identifier of the system on which the job is running and (2) an E if the request is exclusive or an S if the request is shared. RMF reports the jobnames that were active at maximum contention even though those jobs might have been processed and flushed from the system by the time the contention no longer exists. This information can be of great help in balancing your workload to minimize resource contention.

The data gathering technique used to compile the information presented in the report differs from the techniques used for the other reports in that samples are taken only when an enqueue sysevent (ENQHOLD or ENQRLSE) occurs. This technique greatly increases the accuracy of the data, as measurements are taken only when an event occurs that affects the resource contention activity. For more information on the ENQHOLD and ENQRLSE sysevents, see *MVS/XA Initialization and Tuning*, GC28-1149.

The report, in addition to the summary level and the detail 1 level that show lines of data for all resources for which contention occurs, can also be requested for a specific resource by major name, with or without a minor name. Various combinations of the reporting options can give you a complete picture of both critical resources and the jobs that are impacting system throughput by monopolizing a specific resource.

The data fields for the summary and detail reports are identical, with one exception: the job names causing maximum contention are printed only when the detail level is requested. Therefore, the fields are discussed only once, and the field that is provided only at the detail level is noted. The data fields are preceded by ENQUEUE SUMMARY ACTIVITY for a summary report or ENQUEUE DETAIL ACTIVITY for a detail report.

Figure 3-16 shows a sample printed report for enqueue summary activity. Figure 3-17 shows a sample enqueue detail report. Each field is described in Figure 3-18 on page 3-45.

```
                                        E N Q U E U E   A C T I V I T Y
                                                                                                      PAGE 1

             OS/VS2              SYSTEM ID AQFT          DATE 03/04/89          INTERVAL 29.59.997
             SP2.2.3             RPT VERSION 3.5.1       TIME 09.00.00          CYCLE 1.000 SECONDS

ENQUEUE SUMMARY ACTIVITY

-NAME-                                   ----- CONTENTION TIME -----   -%QLEN DISTRIBUTION- AVG Q -REQUEST TYPE - TOTAL
MAJOR                                     MIN    MAX    TOT    AVG      1    2    3    4+  LNGTH -EXCL-  -SHARE- EVENT
 MINOR                                                                                          MIN MAX MIN MAX

HJN
 VERB     (SYSTEMS)
                                          0.027  0.027  0.027  0.027   100  0.0  0.0  0.0   1.00   1   1   0   0

SYSDSN
 D75CEM1.FORM3.LIST(SYSTEMS)
                                          1.327  1.327  1.327  1.327   100  0.0  0.0  0.0   1.00   1   1   0   0

 D10LEM1.PRT.LIST(SYSTEMS)
                                          0.603  0.603  0.603  0.603   100  0.0  0.0  0.0   1.00   1   1   0   0

 BJELLIS.TSOFORM.LIST(SYSTEMS)
                                          0.596  0.596  0.596  0.596   100  0.0  0.0  0.0   1.00   1   1   0   0

SYSIEFSD
 Q10
                                          0.000  0.002  0.033  0.000   98.0  2.0  0.0  0.0  1.02   0   2   0   1    49

 TSOQUE
                                          0.013  1.349  3.320  0.553   100  0.0  0.0  0.0   1.00   1   1   0   0     6

SYSIGGV2
 CATALOG.VTSO025

                                       0.002  0.073  0.134  0.026   100  0.0  0.0  0.0   1.00   0   1   0   1     5

 CATALOG.VCLEAR
                                       0.001  0.309  5.306  0.059    38  .2 61.8  0.0  0.0   1.62   0   2   0   2    89
```

Figure 3-16. Enqueue Activity Summary Report

```
                                        E N Q U E U E   A C T I V I T Y
                                                                                                          PAGE 1

          OS/VS2                      SYSTEM ID AQFT            DATE 03/04/89             INTERVAL 29.59.997
          SP2.2.3                     RPT VERSION 3.5.1         TIME 09.00.00             CYCLE 1.000 SECONDS

ENQUEUE DETAIL ACTIVITY

-NAME-  ----- CONTENTION TIME -----  -- JOBS AT MAXIMUM CONTENTION--  -%QLEN DISTRIBUTION- AVG Q -REQUEST TYPE - TOTAL
MAJOR    MIN    MAX    TOT    AVG    ----- OWN ----- ----- WAIT ----    1    2    3    4+  LNGTH -EXCL-  -SHARE- EVENT
 MINOR                               TOT     NAME    TOT     NAME                               MIN MAX MIN MAX
                                            SYSNAME         SYSNAME
HJN
 VERB    (SYSTEMS)
         0.027  0.027  0.027  0.027    1 #0630092(E)   1 SEGAL    (E)   100  0.0  0.0  0.0   1.00    1   1   0   0      1
                                          AQFT            AQFT
SYSDSN
 D75CEM1.FORM3.LIST(SYSTEMS)
         1.327  1.327  1.327  1.327    1 D75CEM1 (E)   1 D75CEM1Y(E)   100  0.0  0.0  0.0   1.00    1   1   0   0      1
                                          AQFT            AQFT
 D10LEM1.PRT.LIST(SYSTEMS)
         0.603  0.603  0.603  0.603    1 D10LEM1 (E)   1 D10LEM1S(E)   100  0.0  0.0  0.0   1.00    1   1   0   0      1
                                          AQFT            AQFT
 BJELLIS.TSOFORM.LIST(SYSTEMS)
         0.596  0.596  0.596  0.596    1 BJELLIS (E)   1 BJELLISV(E)   100  0.0  0.0  0.0   1.00    1   1   0   0      1
                                          AQFT            AQFT
SYSIEFSD
 Q10
         0.000  0.002  0.033  0.000    1 *MASTER*(S)   1 RCCALL   (E)  98.0  2.0  0.0  0.0   1.02    0   2   0   1     49
                                          AQFT            AQFT
 TSOQUE
         0.013  1.349  3.320  0.553    1 SARAT    (E)  1 GRLEWIS (E)   100  0.0  0.0  0.0   1.00    1   1   0   0      6
                                          AQFT            AQFT
SYSIGGV2
 CATALOG.VTSO025

         0.002  0.073  0.134  0.026    1 CATALOG (E)   1 CATALOG (S)   100  0.0  0.0  0.0   1.00    0   1   0   1      5
                                          AQFT            AQFT
 CATALOG.VCLEAR
         0.001  0.309  5.306  0.059    1 CATALOG (E)   2 CATALOG (E)  38.2 61.8  0.0  0.0   1.62    0   2   0   2     89
                                          AQFT            AQFT
                                                          CATALOG (S)
                                                          AQFT
```

Figure 3-17. Enqueue Activity Detail Report

Figure 3-18. Fields in the Enqueue Activity Report

Heading	Definition
NAME (MAJOR MINOR)	The name of a resource that has one or more requestors waiting for it. The major name is one to eight characters in length; the minor name can be from 1 to 255 characters, but only 44 characters will be printed. When the name exceeds 44 characters, it is truncated in the report, and an asterisk (*) following the resource name indicates that it has been truncated. A resource with a scope of "SYSTEMS" will be followed by (SYSTEMS); a resource with a scope of "SYSTEM" will have no indication; and a resource with a scope of "STEP" will not be included in the report.
CONTENTION TIME (MIN MAX TOT AVG)	The contention time observed for the resource during the RMF reporting interval. The maximum, minimum, total, and average contention times are reported in seconds. The time reported can be 0.000; this indicates a contention time of less than one-thousandth of a second and is most likely to appear as a minimum value. The contention time is calculated by subtracting the time the delay began (when the first ENQHOLD was issued) from the time the contention was ended (when the last ENQRLSE was issued) by the freeing of the resource. An asterisk(*) following the total contention time indicates that the contention extended beyond the end of the measurement interval. The calculation used to determine the average contention time is: `A = T / E` The average contention time in seconds (A) equals the total contention time for the resource during the interval (T) divided by the total number of contention events that occurred during the interval (E).
JOBS AT MAXIMUM CONTENTION	Reported only when the **enqueue activity detail** report is requested, this field describes the total number of resource owners and the total number of jobs waiting to use the resource. In addition, the names of one or two owners and one or two names of waiting jobs are reported. RMF selects the names during the period of maximum contention for each resource. Within this period of maximum contention, RMF determines the point when the queue of waiting jobs was longest and reports the names of the first two jobs on the queue. Each name is followed by an (E) if that job requested exclusive use of the resource or an (S) if that job requested shared use of the resource. Under SYSNAME, RMF reports the name of the system on which the job is executing in a global resource serialization complex. This information can help you to determine which jobs were contributing most heavily to the contention for the resource.
%QLEN DISTRIBUTION (1 2 3 4+)	The percentage of contention events during the interval when the number of requestors queued to the resource was one, two, three, or four or more. The samples are taken during the period when contention was occurring; at each such sample, an accumulator for the observed length is updated. At the end of the measurement interval, the percentage for each queue length is computed. The calculation used for each queue length is: `P = 100 x A/E` The percentage for each queue length (P) equals 100 times the contents of the accumulator for that queue length (A) divided by the total number of enqueue contention events occurring during the interval (E).
AVG Q LNGTH	The average length of the queue of requestors waiting for the resource over the duration of the reporting interval. A consistently high number here indicates that the use of the resource is seriously out of balance. The calculation used is: `A = R / E` The average number of waiting requestors (A) is equal to the total number of requestors waiting during the interval (R) divided by the total number of contention events occurring during the interval (E).
-REQUEST TYPE- -EXCL -- SHARE - (MIN MAX MIN MAX)	The type of the requests — either exclusive or shared — waiting for use of the resource. The requestor would require exclusive use of the resource if the job expects to modify the resource or if the resource is by nature only serially reusable. Other requests would be for shared use of the resource. Both the minimum number and maximum number of waiting shared requests and waiting exclusive requests are reported.
TOTAL EVENT	The total number of resource contention events occurring during the measurement interval.

Messages

During the measurement of enqueue activity, RMF can encounter situations when no reporting can be done. When such a situation occurs, RMF replaces the report with a message describing the reason no report could be formatted. The possible messages are:

NO CONTENTION OCCURRED

Explanation: During the interval, no contention activity occurred for the resource or resources being measured. Enqueue activity measurement and reporting continue as specified. This message would appear most frequently when you are requesting the enqueue activity report for a specific resource.

TABLE FULL-USE SPECIFIC NAME OR SHORT INTERVAL

Explanation: During the RMF interval, a period of such high contention activity occurred that the internal working table was filled. As a result, no further enqueue reporting could be done for the interval. The interval report contains the data gathered before the internal table was filled, followed by the message and subsequent interval reports might not include complete data. Enqueue activity measurement and reporting resume at the start of the next interval.

When the message occurs, you could reduce the length of the RMF interval, or, if you want to ensure that the contention activity for a specific resource is reported, you can request enqueue activity reporting for the specific critical resource. You should also consider enlarging the size of the internal working table. See "RMF Enqueue Contention Buffers" earlier in this chapter.

TERMINATE DUE TO DATA EVENT ERROR - TRY RERUN

Explanation: During the interval, the enqueue measurement routines encountered invalid data while processing a contention event. All enqueue measurement activity is terminated for the session; that is, the recovery from the error includes modifying the enqueue activity option to NOENQ. Because the error encountered might not be a permanent error, you can modify the session options to re-specify enqueue measurement and reporting. If the message occurs again and there are no other indicators of a system problem, report the message to the RMF license holder at your installation.

BAD CPU CLOCK OCCURRED - FIX CPU CLOCK AND RERUN

Explanation: During an interval, the enqueue activity measurement routines detected an error in the CPU clock function. All enqueue measurement activity for the session is terminated; that is, the recovery from the error includes modifying the enqueue activity option to NOENQ. Note that this message is probably one of many indicators that there is a problem with the CPU clock. After the clock has been fixed, re-IPL the system and run the session again, specifying the enqueue activity measurements that you require.

I/O Device Activity Report

The I/O device activity report provides information for all devices in one or more device classes (such as TAPE or DASD) or for those devices you specify on the DEVICE option. When used with the channel path activity and I/O queuing activity reports, this report can help you analyze the I/O activity at your installation and identify bottlenecks caused by a particular device.

Each I/O device activity report begins on a new page, and the class of devices included in the report is indicated by one of the following titles:

```
DIRECT ACCESS DEVICE ACTIVITY
UNIT RECORD DEVICE ACTIVITY
GRAPHICS DEVICE ACTIVITY
COMMUNICATION EQUIPMENT ACTIVITY
MAGNETIC TAPE DEVICE ACTIVITY
CHARACTER READER DEVICE ACTIVITY
```

Note: Because of its speed, the IBM 3890 Document Processor is reported under unit record devices rather than character reader devices.

Because the logical control unit provides a means of identifying a related set of devices, the devices included in the report are grouped by logical control unit. This organization makes it easier to compare the data in the I/O device activity report with the data in the I/O queuing activity report. RMF follows the individual device data lines in each group with a summary line that provides a weighted average or total values for the entire logical control unit.

For the device activity report, the information can be sorted by LCU, or storage group, or both. When the storage group (SG) option is specified, the device activity report is sorted by device number within each storage group. The storage group names are initially sorted by national characters, then alphabetically, and finally, by number. The storage group name that a volume is assigned to is always reported, even when the SG option was not selected. If a volume does not belong to a storage group , the STORAGE GROUP field for that volume is blank. If the volume is added or deleted from a storage group, or the storage management subsystem is not active, the storage group name may not be reported. RMF follows the individual device data lines in each group with a summary line that provides average or total values for the entire storage group.

For any device attached to a byte multiplexor channel, the only measurement data available is the Start Subchannel (SSCH) + Resume Subchannel (RSCH) instruction count. The data line for such a device therefore includes only:

DEV NUM
VOLUME SERIAL
LCU
DEVICE ACTIVITY RATE
%DEV RESV
AVG NUMBER ALLOC
% ANY ALLOC
%MT PEND
%NOT READY

Figure 3-19, Figure 3-20 on page 3-53, and Figure 3-21 on page 3-54 show sample reports for direct access device activity, magnetic tape activity, and communication equipment activity, respectively. The information presented for graphics device activity, unit record device activity, and character reader device activity is the same

 LC28-1556-2

as that shown for communication equipment activity in Figure 3-21. Each data field in this report is described in the following section.

Data lines are included for each device that has been online at least once since Monitor I session initialization. However, data is not reported for devices that were offline at the end of the reporting interval, that came online during the interval, or that were affected by dynamic device reconfiguration during the interval. Such devices are identified under VOLUME SERIAL by one of the following messages in the first data field:

- NOW ONLINE - Brought online during this interval and still online at the end of the interval
- NOW OFFLINE - Taken offline during this interval and still offline at the end of the interval
- OFFLINE - Offline for the entire interval

These messages appear to indicate that the device data is incomplete and might present an inaccurate picture of device activity. Device data might also be incomplete, even when a device has been online for the entire interval, if the hardware measurement data for the device is not available.

If devices are changed or deleted from a storage group, RMF replaces the name of the storage group by **CHGD** in the STORAGE GROUP name column of the direct access device activity report. RMF does not provide summary lines for a storage group with **CHGD** in the STORAGE GROUP name column. Storage group names are still reported when devices are varied on or offline during the interval.

When hardware measurement data is not available, RMF can report values only for fields based on sampled data. It cannot report values for the fields based on hardware measurements; these fields are:

DEVICE ACTIVITY RATE
AVG RESP TIME
AVG DB DELAY
AVG CUB DELAY
AVG IOSQ TIME
AVG PEND TIME
AVG DISC TIME
AVG CONN TIME
%DEV CONN
%DEV UTIL

When it cannot obtain the required data, RMF prints the status message "HARDWARE DATA UNAVAILABLE" in place of the data. Even if the channel measurement facility and the measurement block update facility are active and the device is online for the entire interval, valid hardware data might not be available when the subchannel is in a deferred error state. If RMF is unable to obtain valid hardware data for a subchannel in a deferred error state, it prints the status message "NO H/W DATA."

When RMF cannot calculate an average because a division by zero or a division overflow has occurred, four asterisks (****) appear in the field in place of the data.

Timer Overflow

Depending on the processor model you have, the hardware measurement data might be incomplete because of an overflow in the measurement timer. Because the 308x uses two-byte timers, the maximum connect time, disconnect time, or pending time that can be measured is 8.3 seconds for each I/O request. Any I/O request that exceeds that maximum time limit causes overflow. For example, chain scheduling, which the system uses when accessing page data sets or printing a SYSOUT data set, results in long channel programs and can cause timer overflow. For shared DASD, pending times in excess of 8.3 seconds can occur due to RESERVE activity on the sharing system. Overflow conditions in connect time and disconnect time are detected by the hardware, counted by MVS, and reported by RMF. Overflow conditions in pending time, however, are *not detected.*

The 4381 uses four-byte measurement timers and I/O requests should not exceed the maximum time limit (6.36 days).

Connect/Disconnect Time Overflow

When overflow occurs in connect time or disconnect time, RMF prints "HARDWARE DATA INCOMPLETE" on the line following the requests that caused the overflow. The data presented for those requests is most likely inaccurate because the values shown for connect time and disconnect time per request represent what remained after all the long-running requests were discarded. Those values, along with percent device connected, percent device utilized, and average response time, represent the lower bounds of what the actual values might be. Because at least one request was discarded, all values must be larger than reported; how much larger, however, cannot be exactly determined.

On the same line with "HARDWARE DATA INCOMPLETE," RMF records the values of two counters: total requests that had timer overflow (in either or both timers); and total requests that had connect time overflow. The difference between timer overflow and connect time overflow gives the number of requests that had only disconnect time overflow. If the difference is large, it might indicate that difficulty in reconnecting to the channel is causing delays. Total requests with timer overflow is the difference between the start subchannel count and the measurement event count. RMF records connect time overflow separately.

The counts of requests that had overflow are intended to indicate how much data was lost. For example, if 1000 requests occurred in an hour and only one had overflow, the actual values are probably not much larger than the reported values. However, if the 999 requests measured were all short and all occurred within a short span of time, whereas, the one long request lasted for 95% of the hour, the reported data is highly inaccurate. The values reported in the two counters do not take into account how many times a single request had an overflow.

The report of overflows for paging devices does not necessarily indicate a problem. Consult the page/swap data set report and the workload activity report to determine whether or not paging delays are a problem. If they are, the device data can be used in conjunction with the two reports to analyze the problem.

Pending Time Overflow

Because all overflows for pending time are lost, RMF does not accurately report certain shared DASD delays. For example, a request delayed for 18 seconds overflows twice; 16.6 seconds are lost. To RMF, the delay appears to be only 1.4 seconds. Therefore, the AVERAGE PENDING TIME and the AVERAGE RESPONSE TIME values are extremely inaccurate.

For requests with extremely long delays, the missing interrupt handler (MIH) halts the request and reschedules it periodically. MIH estimates the amount of pending time, based on the MIH interval, and adds it to the value RMF reports. Therefore, pending time is lost only for requests that take longer than 8.3 seconds and less than 1.5 times the MIH interval. To increase the accuracy of AVERAGE PENDING TIME and AVERAGE RESPONSE TIME, decrease the MIH interval. An interval of 4 seconds will ensure that no pending time is lost. However, some performance penalty does occur because of the 4-second interval.

Other Reports

In Monitor II reports, the overflow counts are not reported. A value affected by overflow, however, is marked by an asterisk (*). In exception reports, plot reports and summary reports, overflows are neither detected nor identified. Because Monitor II device activity reports, as well as some exception reports and plot reports and some fields in the summary report, are based on data the Monitor I session collects, the same inaccuracies apply to these reports.

Duration Report

When you request a duration report for I/O device activity, the identifier of each I/O device that had any VARY activity during the duration interval is followed by a single asterisk (*). The data recorded for such a device is partial; that is, no data was collected during one or more of the measurement intervals included in the duration report because the device was varied during a measurement interval.

The calculations the post processor performs to generate a duration report make no adjustments for RMF measurement intervals when no data was collected for a device. Thus, the data for a device that moved online or offline might appear to be inconsistent. For example, the percentages reported for the QLENGTH DISTRIBUTION field do not equal 100% when data for the device is partial.

At the beginning of a Monitor I session, all devices that are online are known to RMF; thus, RMF creates an entry in the type 74 SMF record for each online device that the user requested RMF to monitor. In contrast, any device that is offline at the beginning of the session is unknown to RMF, and no entry in the SMF record is built. When an unknown device is brought online, it becomes known to RMF, and an entry in the SMF record is then built for the device.

When a duration report combines data collected during two or more separate Monitor I sessions, the status of a device can change (for example, from offline to online or from unknown to online). Thus, the following conditions can occur:

- A device was known to RMF but offline during the first Monitor I session and online during subsequent Monitor I sessions but did not change during a measurement interval included in the duration report. In this case, the data is partial and the device identifier is followed by an asterisk.
- A device was offline for one or more measurement intervals and unknown to RMF during all other measurement intervals included in the duration report. In this case, the device identifier is followed by an asterisk, OFFLINE appears in the first data field, and no data is formatted for the device.

- A device was unknown to RMF during a Monitor I session and online for subsequent Monitor I sessions but did not change during a measurement interval included in the duration report. In this case, the data is partial and the device identifier is followed by an asterisk.

The following conditions can occur for storage group reporting:

- The STORAGE GROUP field shows **CHGD** for the volume if the storage group name changes in an SMF record for the duration period.
- The STORAGE GROUP field contains a valid storage group name and the LCU field contains an asterisk if a volume is not reported in all SMF records of that duration period, but the storage group name is the same in all SMF records that are part of the duration report.
- The STORAGE GROUP field shows **CHGD** and the LCU field contains an asterisk if a volume is not reported in all SMF records of the duration and has changed the storage group name at least once.
- STORAGE GROUP DATA NOT AVAILABLE is reported between the TOTAL SAMPLES = field and the report headings if the storage management subsystem is not available in one of the reports during the duration period.
- SMS INTERFACE ERROR, NEW STORAGE GROUP INFORMATION CANNOT BE OBTAINED is reported between the TOTAL SAMPLES = field and the report headings if a Storage Management Subsystem interface error occurs in one of the reports during the duration period.

D I R E C T   A C C E S S   D E V I C E   A C T I V I T Y

PAGE 1

OS/VS2 SYSTEM ID 168A DATE 04/15/89 INTERVAL 01.33.576
SP2.2.3 RPT VERSION 3.5.1 TIME 10.34.27 CYCLE 1.000 SECONDS

TOTAL SAMPLES = 60

STORAGE GROUP	DEV NUM	VOLUME SERIAL	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG CUB DELAY	AVG DB DELAY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NUMBER ALLOC	% ANY ALLOC	% MT PEND
	220	DRV021		0.011	11	0			0.1	4.4	6.7	0.01	0.01	0.0	11.0	100.0	0.0
	270	PAGE08		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	6.0	100.0	0.0
	271	SPOOL1		0.321	24	0			0.2	20.6	3.5	0.11	0.77	0.0	6.0	100.0	0.0
	272	RMFPAK		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
	273	SMSDAT		0.085	11	0			0.2	7.4	3.8	0.03	0.10	0.0	3.0	100.0	0.0
MVSXA	274	GNRL01		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
VIRTUAL	275	GNRL02		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
COMMON	276	GNRL03		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
COMMON	277	TSOX01		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
	278	RMFSYS		0.032	16	0			1.4	4.7	10.2	0.03	0.05	0.0	1.0	100.0	0.0
	291	191MJK		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
COMMON	2C0	DUMY01		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
COMMON	2C1	ALL001		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
EXCLUS	2C2	PUB001		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
ALLGRP	2C3	SEC001		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
	2C9	RMFDMP		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0
	2CA	RMFSMF		0.000	0	0			0.0	0.0	0.0	0.00	0.00	0.0	0.0	100.0	0.0

Figure 3-19. Direct Access Device Activity Report

M A G N E T I C   T A P E   D E V I C E   A C T I V I T Y

PAGE 1

OS/VS2 SYSTEM ID AQFT DATE 03/04/89 INTERVAL 29.59.997
SP2.2.3 RPT VERSION 3.5.1 TIME 09.00.00 CYCLE 1.000 SECONDS

TOTAL SAMPLES = 1,800

DEV NUM	VOLUME SERIAL	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG CUB DELAY	AVG DB DELAY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NUMBER ALLOC	% ANY ALLOC	% MT PEND	% NOT READY
580		031	1.041	928	880	0.3	0.0	28.4	0.1	20.1	2.09	2.10	0.0	0.1	9.0	0.6	91.6
581		031	0.000	****	****	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	0.0	0.0	0.0	100.0
582	T80686	031	4.398	225	192	0.0	0.0	15.3	0.0	18.3	8.04	8.05	0.0	1.0	100.0	84.3	84.3
588		031	1.109	864	804	0.1	0.0	41.3	0.0	18.4	2.04	2.04	0.0	0.1	10.8	0.0	89.2
589		031	1.148	832	771	0.2	0.0	42.6	0.1	18.3	2.10	2.11	0.0	0.1	11.5	2.5	91.0
58A		031	1.065	549	501	0.0	0.0	27.9	0.1	19.7	2.10	2.11	0.0	0.1	22.2	8.6	86.4
	LCU	031	8.761	****	****	0.1	0.0	25.3	0.0	18.7	2.73	2.73	0.0	1.4	25.6	16.0	90.4
5A0	T0171	033	5.440	77	64	0.1	0.0	1.6	2.9	8.5	4.63	6.21	0.0	0.4	42.1	34.7	92.6
5A1		033	0.000	****	****	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	0.0	0.0	0.0	100.0
5A2		033	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	0.0	0.0	0.0	100.0
5A3		033	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	0.0	0.0	0.0	100.0
5A4		033	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	0.0	0.0	0.0	100.0
5A5		033	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	0.0	0.0	0.0	100.0
5A6		033	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	0.0	0.0	0.0	100.0
5A7		033	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0	0.0	0.0	0.0	100.0
	LCU	033	5.440	****	****	0.1	0.0	1.6	2.9	8.5	0.58	0.78	0.0	0.4	5.3	4.3	99.1

Figure 3-20. Magnetic Tape Device Activity Report

C O M M U N I C A T I O N   E Q U I P M E N T   A C T I V I T Y

PAGE 1

OS/VS2 SP2.2.3 — SYSTEM ID AQFT, RPT VERSION 3.5.1 — DATE 03/04/89, TIME 09.00.00 — INTERVAL 29.59.997, CYCLE 1.000 SECONDS

TOTAL SAMPLES = 1,800

DEV NUM	VOLUME SERIAL	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG CUB DELAY	AVG DB DELAY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NUMBER ALLOC	% ANY ALLOC	% MT PEND	% NOT READY
C44		083	91.015	0	0	0.0	0.0	0.1	0.0	0.2	2.26	2.69	0.0		100.0		0.0
C45		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C46		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C47		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C4C		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C4D		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C4E		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C4F		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C50		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C51		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C52		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C53		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C54		083	45.507	2	0	0.0	0.0	0.1	1.7	0.6	2.80	10.74	0.0		100.0		0.0
C55		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C56		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C57		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C5D		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
C5F		083	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
	LCU	083	136.521	1	0	0.0	0.0	0.1	0.6	0.4	0.28	0.75	0.0		11.1		0.0
CC0		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CC1		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CC2		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CC3		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CC4		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CC5		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CC6		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CC7		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CC8		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CC9		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CCA		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CCB		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CCC		084	0.693	106	0	0.0	0.0	0.1	74.0	31.4	2.18	7.30	0.0		100.0		0.0
CCD		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CCE		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CCF		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CD0		084	0.205	58	0	0.0	0.0	0.2	26.5	31.4	0.64	1.19	0.0		100.0		0.0
CD1		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CD2		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0
CD3		084	0.000	0	0	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.0		0.0		0.0

Figure 3-21. Communication Equipment Activity Report

Figure 3-22 (Page 1 of 5). Fields in the I/O Device Activity Reports

Heading	Definition
STORAGE GROUP	The name that identifies the storage group to which the device belongs. Your storage administrator assigns the names. These names are available on the direct access device report only.
DEV NUM	A three-digit hexadecimal device number that identifies a physical I/O device.
VOLUME SERIAL	The volume serial number (for direct access and magnetic tape reports) of the volume mounted on the device at the end of the reporting interval.
LCU	A number of the logical control unit to which the device belongs. The definition of a logical control unit (LCU) is model-dependent. On a 308x or 908x processor, an LCU is the set of devices attached to the same physical control unit (or a group of physical control units with one or more devices in common.) Any I/O request for any device attached to any physical control unit in the group is queued at the logical control unit for the 308x or 908x processors. For a 3090 processor, the definition of an LCU is the same as that for a 308x processor, but the IOP, which is part of the channel subsystem, manages and schedules I/O work requests. In any event, to find the logical control unit number for the 308x, 908x, and 3090 processors, RMF must access the I/O configuration data set (IOCDS). If RMF cannot read it, or if it has been updated so that the data might not apply to the present configuration, RMF issues a status message, IOCDS INFORMATION UNAVAILABLE TO RMF. Two different response codes are issued: A response code of 2024 indicates that the IOCDS has been updated so that the data might not apply to the present configuration. A response code of 43F0 indicates that RMF cannot read the IOCDS because it appears in another partition of a multi- processing system (for example, when RMF is running on a 3084 processor, and the operator partitions the system in such a way that RMF can no longer read the IOCDS). In either case, RMF cannot obtain the LCU identifier type. RMF continues to measure I/O device activity, but the requested SMF record does not contain the LCU identifier. The LCU field remains blank, and the device activity report lists the device numbers in ascending order. On a 4381 processor, an LCU is the set of devices attached to an identical set of physical control units. While I/O request measurements for the set of devices in an LCU are maintained on a physical control unit basis, they are summed together by the channel microcode when RMF retrieves these measurements on an LCU basis. In this way, RMF provides one set measurements for the group of devices in an LCU. To find the logical control unit number, RMF uses the I/O configuration data (IOCD) in the hardware system area (HSA). If RMF cannot retrieve the configuration data after successive attempts, RMF does not report the logical control unit number and does not summarize device data by LCU.
DEVICE ACTIVITY RATE	The rate per second at which start subchannel (SSCH) instructions to the device completed successfully. In the line of data for the base exposure of a multiple exposure device, this field reflects the rate for the entire device. The data line for the base exposure is followed by data lines for each exposure. The value is calculated by dividing the total number of successful SSCH instructions by the number of seconds in the interval. For devices using suspended channel programs, resume I/O requests are included in the SSCH counts. In the logical control unit summary line, this field contains the sum of the rates for each individual device. If there is a multiple exposure device belonging to the logical control unit, only the value represented by the base exposure is considered in the summary line.

Figure 3-22 (Page 2 of 5). Fields in the I/O Device Activity Reports

Heading	Definition
AVG RESP TIME	The average number of milliseconds the device required to complete an I/O request. This value reflects the total hardware service time and the front end software queuing time involved for the average I/O request to the device. The channel measures active time, which starts at the acceptance of a SSCH instruction (indicated by a condition code 0) and ends at the acceptance of the channel end (primary status pending). It does not, however, include the time required to process the interruption. The IOS queue length is factored in to reflect the front end queuing time. The calculation used is: `ART = (AT/MEC) + QT` The average response time (ART) equals the active time (AT) for the device divided by the measurement event count (MEC), plus the average time a request waited in the subchannel queue (QT). (RMF reports this average queuing delay under AVG IOSQ TIME.) NOTE: Active time is the sum of connect time, disconnect time, and pending time. This field is the sum of AVG IOSQ TIME, AVG PEND TIME, AVG DISC TIME, and AVG CONN TIME described later. For the base exposure line of a multiple exposure device, this field contains the weighted average of response times for the non-base exposures. In the logical control unit summary line, this field contains the weighted average of the individual average response times for each device.
AVG IOSQ TIME	The average number of milliseconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued. Delay occurs when a previous request to the same subchannel is in progress. The calculation is: `AQT = (QC/S) / R` The average IOS queue time (AQT) equals the IOS queue count (QC) divided by the number of samples (S), divided by the device activity rate (R). (RMF calculates and reports this rate under DEVICE ACTIVITY RATE.) In the logical control unit summary line, this field contains the weighted average of the individual average queue times of each device.
AVG CUB DELAY	The average number of milliseconds of delay an I/O request encountered because the control unit was busy. The data field heading appears in the report for all processors, but data appears in the field only when RMF is running on a 3090 processor. Otherwise, the field is left blank. If the device is shared at the control unit level, the contention might be caused by the sharing system. If the device is not shared at the control unit level, the contention is the result of other activity to different devices over the same or an alternate path serviced by the control unit. The calculation used is: `ACUB = CUBD / MEC` The average control unit busy delay time (ACUB) equals the control unit busy delay time (CUBD) divided by the measurement event count (MEC). In the line of data for the base exposure of a multiple exposure device, this field is the weighted average of the corresponding values for all exposures. In the LCU summary line, this field contains the weighted average of the individual AVG CUB DELAY times.

Figure 3-22 (Page 3 of 5). Fields in the I/O Device Activity Reports

Heading	Definition
AVG DB DELAY	The average number of milliseconds of delay that I/O requests to this device encountered because the device was busy. The data field heading appears in the reports for all processors, but data appears in the field only when RMF is running on a 3090 processor. Otherwise, the field is left blank. Device busy might mean: 1. The volume is in use by another system. 2. The device is reserved by another system. 3. Contention was caused by head of string busy conditions. 4. Some combination of these three conditions has occurred. The calculation used is: `ADB = DBD / MEC` The average device busy delay time (ADB) equals the device busy delay time (DBD) divided by the measurement event count (MEC). In the line of data for the base exposure of a multiple exposure device, this field is the weighted average of the corresponding values for all exposures. In the LCU summary line, this field contains the weighted average of the individual AVG DB DELAY times.
AVG PEND TIME	The average number of milliseconds an I/O request must wait in the hardware. This value reflects the time between acceptance of the SSCH function at the channel (SSCH-function pending) and acceptance of the first command associated with the SSCH function at the device (subchannel active). This value also includes the time waiting for an available channel path and control unit as well as the delay due to shared DASD contention. If the value is high, refer to the device's LCU entry in the I/O queuing activity report for an indicator of the major cause of the delay. The calculation is: `APT = FPT / MEC` The average pending time (APT) equals the total function pending time for the device (FPT) during the measurement interval divided by the measurement event count (MEC). In the logical control unit summary line, this field contains the weighted average of the individual average pending times for each device. For the base exposure line of a multiple exposure device, this field contains the weighted average of the average pending time for the non-base exposures.
AVG DISC TIME	The average number of milliseconds the device was disconnected while processing an SSCH instruction. Thus, this value reflects the time when the device was in use but not transferring data. It includes the overhead time when a device might disconnect to perform positioning functions such as SEEK/SET SECTOR, as well as any reconnection delay. The calculation is: `ADT = TDT / MEC` The average disconnect time (ADT) equals the total disconnect time for the device (TDT) during the interval divided by the measurement event count (MEC). The MEC is the same as the number of SSCH instructions issued, unless there has been a timer overflow error in the channel. For the base exposure line of a multiple exposure device, this field contains the weighted average of the average disconnect time for the non-base exposures. In the logical control unit summary line, this field contains the weighted average of the individual average disconnect times for each device.

Figure 3-22 (Page 4 of 5). Fields in the I/O Device Activity Reports	
Heading	**Definition**
AVG CONN TIME	The average number of milliseconds the device was connected to a channel path and actually transferring data between the device and main storage. Typically, this value, measures data transfer time but also includes the search time needed to maintain channel path, control unit, and device connection. The calculation used is: `ACT = DCT / MEC` The average connect time (ACT) equals the total connect time for the device (DCT) during the measurement interval divided by the measurement event count (MEC). For the base exposure line of a multiple exposure device, this field contains the weighted average of the average connect time for the non-base exposures. In the logical control unit summary line, this data field contains the weighted average of the individual average connect times for each device.
% DEV CONN	The percentage of time during the interval when the device was connected to a channel path. The calculation used is: `DCP = 100 x (DCT/T)` The device connect percentage (DCP) equals 100 multiplied by the total device connect time (DCT) divided by the length of the measurement interval (T). For the base exposure line of a multiple exposure device, this field contains the sum of the corresponding values for all exposures. In the logical control unit summary line, this field contains the average of the individual device connect percentages for each device.
% DEV UTIL	The percentage of time during the interval when the device was in use. This percentage includes both the time when the device was involved in I/O operations (connect and disconnect time) and the time when it was reserved but not involved in an I/O operation. The percentage reported represents the time during the interval when the device is "tied up" when it could not be used to service a request from another system. RMF uses hardware measurement values to compute the time when the device was involved in an I/O operation. It uses sampling techniques to estimate the time when the device was reserved but not actually involved in an I/O operation, based on the settings of two bits: UCBSTRT — is set on when the device is processing an I/O request and off when there is no active I/O request. UCBRESVH — is on when the device is reserved. When UCBSTRT is off and UCBRESVH is on, RMF counts the sample as indicating non-overlapped device reserved time; that is, the device is reserved but not actively performing an I/O operation. Some small portion of device busy (reserved) time is missed when the device is reserved and UCBSTRT is on but the I/O request is pending in the channel. The calculation used is: `DUP = 100 x (DCT+TDT) / T + (RS/S)` The device utilization percentage (DUP) equals 100 times the sum of the device connect and disconnect time (DCT + TDT), divided by the length of the measurement interval in seconds (T) and the number of samples when the device was reserved but not involved in an I/O operation (RS) divided by the total number of samples (S). For the base exposure line of a multiple exposure device, this field contains the weighted average of the corresponding values for the non-base exposures. In the logical control unit summary line, this data field contains the average of the individual percentages reported for each device.

Figure 3-22 (Page 5 of 5). Fields in the I/O Device Activity Reports

Heading	Definition
% DEV RESV	The percentage of time during the interval when a shared device was reserved by the processor on which RMF was started. At each RMF cycle, RMF checks to see if a device is reserved, and a counter is kept of all such samples. At the end of the interval, the percentage is computed. The calculation used is: `DRP = 100 x DRS/S` The percentage of time during the interval when the device is reserved (DRP) equals 100 times the number of device-reserved samples (DRS) divided by the total number of samples (S). In the logical control unit summary line, this field contains the average of the individual device reserved percentages for each device.
AVG NUMBER ALLOC	The average number of data control blocks (DCBs) and access method control blocks (ACBs) concurrently allocated for each volume. This field is reported only for direct access storage devices and magnetic tape devices. At each RMF cycle, a counter is increased to reflect the number of data sets concurrently allocated. At the end of the interval, the average is calculated by dividing the total number of allocated data sets for all samples by the total number of samples. In the logical control unit summary line, this field contains the average of the individual average (S) number of allocated data sets for each device.
ANY ALLOC	The percentage of time during the reporting interval when the device was allocated to one or more data sets. Permanently mounted direct access devices show a 100% allocation, regardless of whether or not a data set was actually allocated. To determine the value, RMF keeps a count at each cycle of whether or not the device was allocated or permanently resident. At the end of the interval, the percentage is computed. The calculation used is: `P = 100 x AS/S` The percentage (P) equals 100 times the number of allocated samples (AS) divided by the total number of samples (S). In the logical control unit summary line, this field contains the average of the individual allocation percentages for each device.
% MT PEND	The percentage of time during the interval when a mount was pending for the device. This field is reported only for direct access devices and magnetic tape devices. At each cycle, RMF updates a counter when it detects a mount pending condition. At the end of the interval, the percentage is computed. The calculation used is: `MPP = 100 x C/S` The mount pending percentage (MPP) equals 100 times the value in the counter (C), divided by the total number of samples (S). In the logical control unit summary line, this field contains the average of the individual mount pending percentages for each device.
% NOT READY	The percentage of time during the reporting interval when the device was not ready for use, as, for example, when a tape has just been mounted but is not yet ready to be used to the system. This field is not reported for direct access devices; however, the value is recorded in the corresponding field of the SMF record, should your installation need the information. At each RMF cycle, a counter is updated when the status of the device indicates that it is not ready. At the end of the interval, the percentage is computed. The calculation used is: `NRP = 100 x NRS/S` The not ready percentage (NRP) equals 100 times the count of not ready samples (NRS), divided by the total number of samples (S). In the logical control unit summary line, this field contains the average of the individual not ready percentages for each device.

I/O Queuing Activity Report

The I/O queuing activity report provides information on the I/O configuration and activity rate, queue lengths, and percentages when one or more I/O components were busy, grouped by logical control unit. The format and contents of the report, like the definition of an LCU, are model-dependent. The following three sections explain the differences in the report for the model processor that you use.

I/O Queuing Activity Report for 308X, 908X, or 4381 Processors

Figure 3-23 gives the format of the report, and Figure 3-24 describes the content and meaning of each field. On a 308x or 908x processor, an LCU is the set of devices attached to the same physical control unit (or a group of physical control units with one or more devices in common.) Each physical control unit and each device belong to only one LCU, and all requests to all devices in the LCU are placed on the single queue for that unit, regardless of whether or not there is contention. (Although IOS starts an I/O request to a subchannel, the request is actually queued by the LCU.)

Note that for the 4381 processor, the AVG Q LNGTH field heading appears in the report but the field is left blank. On a 4381 processor, an LCU is the set of devices attached to an identical set of physical control units. A physical control unit can belong to more than one LCU if the devices are configured asymmetrically. Each device belongs to only one LCU. I/O request initiation is performed on a physical control unit basis.

Your installation defines your I/O configuration as input to the input/output configuration program (IOCP). The IOCP uses the information you supply to define the relationship between channel paths, control units, and I/O devices. The IOCP generates and assigns LCU identifiers to these groups of channel paths, control units and I/O devices. The IOCP then places this configuration definition in a configuration data set (IOCDS) for the 308x or 908X processors. The IOCP places this configuration definition in the input/output configuration data (IOCD) in the hardware system area (HSA) in a 4381 processor complex. RMF uses the configuration definition as well as measurement data gathered during the interval to generate the I/O queuing activity report.

A channel path can be online to the system but not connected to any device in an LCU during the measurement interval. In this case, a line identifies that channel path as PATH(S) OFFLINE. If an online channel path was not connected to any device of the LCU at the last interval but is now connected to one or more devices of an LCU because of vary path processing, a line identifies that channel path PATH(S) NOW ONLINE.

I/O Queuing Activity Report for 3090 Processors

Figure 3-25 shows the format of the report, and Figure 3-26 describes the content and meaning of each field.

On a 3090 processor, an LCU is the set of devices attached to the same physical control unit (or group of control units that have one or more devices in common). As with the 308x processor, each device belongs to only one LCU, but the I/O processor (IOP), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU of a 3090 processor. If an I/O request is unsuccessful because the control unit is busy, the request is queued on the control unit header (CU-HDR) queue. Once the busy condition is resolved, the CU-HDR is then placed on the initiative queue.

Your installation defines your I/O configuration as input to the input/output configuration program (IOCP). The IOCP uses the information you supply to define the relationship between channel paths, control units, and I/O devices. The IOCP generates and assigns LCU identifiers to these groups of channel paths, control units and I/O devices. The IOCP then places this configuration definition in a configuration data set (IOCDS) in a 3090 processor. RMF uses the configuration definition as well as measurement data gathered during the interval to generate the I/O queuing activity report.

For all processors

The information in the I/O queuing activity report, used after the channel path activity and I/O device activity reports have shown that a problem exists, can help you pinpoint the reason for contention delays associated with channel paths, control units, and devices. For example, if the I/O device activity report shows an unusually large pending time for one or more devices in a logical control unit, the I/O queuing activity report indicates what proportion of the delay is caused by control unit busy and device busy, thus indicating which part of the configuration might need adjustment.

You can also use the two reports to analyze the current I/O configuration. The I/O device activity report shows which devices belong to each logical control unit. The I/O queuing activity report shows which physical control units are part of each logical control unit and which channel paths are connected to each physical control unit.

When a logical control unit has no activity during the interval, RMF omits that logical control unit from the report for that interval. If no activity has occurred during the interval for all selected LCUs, the message NO ACTIVITY FOR SELECTED LCUs appears instead of the data after the headings of the report.

If a channel path was brought online or taken offline during the interval, data is formatted and an additional line in the report describes its status. If an installed channel path was offline during the whole interval, the additional line identifies the channel path as OFFLINE. If a channel path was taken offline or brought online during an interval, the additional line identifies the channel path as either NOW OFFLINE or NOW ONLINE.

When RMF is unable to obtain valid hardware data for CONTENTION RATE and DELAY Q LNGTH in the 3090 I/O queuing activity report, it prints the message NO H/W DATA under those headings. If the channel measurement facility is inactive or has been interrupted during the interval, CHANNEL MEASUREMENT FACILITY NOT ACTIVE OR INTERRUPTED appears after the headings where the data normally appears in the 3090 I/O queuing activity report. If the diagnosis interface fails during the interval, DIAGNOSIS INTERFACE FAILURE appears after the headings in the 3090 I/O queuing activity report.

The reports depend on information in the I/O configuration data set (IOCDS). If RMF cannot read the IOCDS, or if the IOCDS has been updated so that the data might not apply to the present configuration, no report is available.

For example, when RMF is running on a 3084 processor and the operator partitions the system in such a way that RMF cannot read the IOCDS because it appears in another partition of a multi-processing system, RMF terminates the I/O queuing activity report and issues to the operator a message, I/O QUEUING ACTIVITY RMF REPORT TERMINATED.

```
                                       I/O  Q U E U I N G  A C T I V I T Y
                                                                                         PAGE 1

          OS/VS2                   SYSTEM ID AQFT            DATE 03/04/89            INTERVAL 29.59.997
          SP2.2.3                  RPT VERSION 3.5.1         TIME 09.00.00            CYCLE 1.000 SECONDS

  TOTAL SAMPLES =  1800


                                  % ALL                --% REQ DEFER --
LCU       ACTIVITY  AVG Q   CH PATH     % REQ     DEV      CU          CONTROL      CHAN
          RATE      LNGTH   BUSY        DEFER     BUSY     BUSY        UNITS        PATHS

000       --------  -----   0.00        -----     ----     ----        009          10,50 DATA INVALID STSCH REQUEST FAILED
005          1.915   0.00   0.11         0.00      0.00     0.00       007          10,50
                                                                                    50    PATH(S) OFFLINE
00A         16.693   0.53   8.11        15.37      0.00    15.37       010          01,41
                                                                       011          11,51
00B         17.230   0.38   8.11        14.32      0.00    14.32       018          01,41
00B         17.230   0.38   8.11        14.32      0.00    14.32       018          01,41
00B         17.230   0.38   8.11        14.32      0.00    14.32       018          01,41
00B         17.230   0.38   8.11        14.32      0.00    14.32       018          01,41
00B         17.230   0.38   8.11        14.32      0.00    14.32       018          01,41
00B         17.230   0.38   8.11        14.32      0.00    14.32       018          01,41
00B         17.230   0.38   8.11        14.32      0.00    14.32       018          01,41
00B         17.230   0.38   8.11        14.32      0.00    14.32       018          01,41
```

Figure 3-23. I/O Queuing Activity Interval Report (308x, 908x, 4381)

Figure 3-24 (Page 1 of 2). Fields in the I/O Queuing Activity Report (308x, 908x, 4381)	
Field Heading	**Definition**
LCU	The three-digit hexadecimal number that identifies the logical control unit (LCU). The set of devices associated with an LCU measurement are not the same on all processors because the definition of an LCU is model-dependent. **On a 308x or 908x processor complex**, an LCU is the logical representation of a physical control unit or a group of physical control units with one or more devices in common. Each physical control unit and each device can belong to only one LCU; they cannot be shared between LCUs. To find the LCU number, RMF must access the I/O configuration data set (IOCDS). If RMF cannot read it, or if it has been updated so that the data might not apply to the present configuration, RMF ends I/O queuing activity and does not produce a report. If no activity has occurred during the interval for all selected LCUs, the message NO ACTIVITY FOR SELECTED LCUs appears instead of the data after the headings of the report. **On a 4381 processor complex**, an LCU is the logical representation of a set of devices attached to an identical set of physical control units. Each device belongs to only one LCU, however, a physical control unit may belong to more than one LCU. To find the LCU number, RMF must access the I/O configuration data (IOCD) in the hardware system area (HSA) in a 4381 processor. If RMF cannot read it, or if it has been updated so that the data might not apply to the present configuration, RMF ends I/O queuing activity and does not produce a report. If no activity has occurred during the interval for all selected LCUs, the message NO ACTIVITY FOR SELECTED LCUs appears instead of the data after the headings of the report.
ACTIVITY RATE	The rate per second at which I/O requests were successfully selected for initiation. When this rate is relatively high, the control units in the LCU might be overloaded. The activity rate is calculated by dividing the total number of I/O requests successfully selected during the interval by the number of seconds in the interval.
AVG Q LNGTH	**For a 308x/908x processor complex**, this field appears when I/O queuing activity is measured and is the average number of requests on the logical control unit (LCU) queue. Each time the channel subsystem takes a request off the queue it counts the number of requests remaining and adds this number to an accumulator. At the end of the interval, RMF divides the accumulated number of requests by the number of times an I/O request was dequeued. The calculation used is: `AQL = ER / DR` The average queue length (AQL) for the interval equals the total number of enqueued requests (ER) divided by the number of times an I/O request was dequeued (DR). **For a 4381 processor complex**, I/O requests are queued on a physical control unit basis within the channel subsystem. Because queue length information is not available from the channel subsystem, the AVG Q LNGTH field is blank in the report for a 4381 processor, although the field heading does appear.
% ALL CH PATH BUSY	The percentage of time during the measurement interval when all channel paths belonging to the logical control unit were busy at the same time. At each sampling cycle, RMF checks the results of the last Store Channel Path Status (STCPS) instruction that the SRM issued to determine whether all channel paths are busy. RMF performs this check only for those channel paths that are online at that sample. At the end of the interval, RMF computes the percentage. The calculation used is: `ACB = 100 x (SB/S)` The percentage of time during the interval when all channel paths in the logical control unit were busy (ACB) equals 100 times the number of samples that indicated all channel paths were busy (SB) divided by the number of samples taken (S).

Figure 3-24 (Page 2 of 2). Fields in the I/O Queuing Activity Report (308x, 908x, 4381)

Field Heading	Definition
% REQ DEFER	The percentage of unsuccessful attempts to start I/O requests during the measurement interval. Requests are deferred mainly because the control unit is busy or the device is busy. Other reasons for delay might be a disabled path. The calculation used is: `RDP = 100 x (RD/RA)` The percentage of requests deferred (RDP) equals 100 times the number of requests deferred (RD) divided by the number of attempts the channel made to initiate I/O requests (RA). Note that RMF calculates the number of requests deferred (RD) by subtracting the successful attempts from the total. **Note:** The percentage of deferred requests may be quite high under normal conditions because many of the requests counted are not actually delayed. For example, if a device has two control units and one is busy, the I/O request may be "deferred" before selection is successful. However, the request is only delayed for the time it took the channel to find the free path, which is insignificant.

% REQ DEFER DEV BSY CU BSY	The percentage of I/O requests deferred because the device was busy (DEV BUSY) or because the control unit was busy (CU BUSY). RMF uses a separate calculation for each percentage.
	
	To determine the percentage of requests deferred because the device was busy, the calculation is:
	
	`PDB = 100 x (RDB/RA)`
	
	The percentage (PDB) equals 100 times the number of requests deferred because the device was busy (RDB) divided by the total number of attempts the channel made to initiate I/O requests (RA).
	
	To determine the percentage of requests deferred because the control unit was busy, the calculation is:
	
	`PCB = 100 x (RCB/RA)`
	
	The percentage (PCB) equals 100 times the number of requests deferred because the control unit was busy (RCB) divided by the total number of attempts the channel made to initiate I/O requests (RA).
	
	Note: The absolute size of the number of requests deferred is really insignificant (see the note under "% REQ DEFER"). The relative proportion of requests deferred due to device unit busy and control unit busy can be used to determine the most probable cause of delays after you use the other reports to determine if delays exist.

CONTROL UNITS	The three-digit hexadecimal identifier of each physical control unit contained in the logical control unit. There can be a maximum of four physical control units associated with a logical control unit. On a 4381 processor, physical control units can be shared among logical control units.

CHAN PATHS	The two-digit hexadecimal channel path identifier (CHPID) of the channel paths that are attached to the physical control units contained in the logical control unit. There can be up to four channel paths in a logical control unit. All channel paths are listed, even those that are offline. The channel paths that are offline or changed online or offline during the interval are indicated as follows:
	
	`'OFFLINE'`
	`'NOW OFFLINE'`
	`'NOW ONLINE'`
	
	Channel paths that are online to the system but that might or might not be connected during the interval to any device in an LCU are indicated as follows:
	
	`'PATH(S) OFFLINE'`
	`'PATH(S) NOW OFFLINE'`
	`'PATH(S) NOW ONLINE'`


```
                                   I/O  Q U E U I N G   A C T I V I T Y
                                                                                     PAGE 1

        OS/VS2                  SYSTEM ID AQFT         DATE 03/04/89           INTERVAL 29.59.997
        SP2.2.3                 RPT VERSION 3.5.1      TIME 09.00.00           CYCLE 1.000 SECONDS

TOTAL SAMPLES =  1800      IOP   ACTIVITY RATE  AVG Q LNGTH
                           00      315.974         0.01
                           01      553.859         0.01

                    DELAY    % ALL
LCU  CONTENTION      Q      CH PATH     CONTROL    CHAN      CHPID    % CU
        RATE       LNGTH     BUSY        UNITS     PATHS     TAKEN    BUSY
012     0.069       0.06     0.00         020       1B      18.031   14.46
                                                    5B      20.448   10.10
                                          021       2F      20.944    7.13
                                                    65      16.883   18.03
015     0.002       0.20     0.00         028       1B       3.001    2.77
                                                    5B       3.078    2.03
                                          029       2F       3.018    1.54
                                                    65       2.927    3.04
017     0.000       0.00     0.01         022       1B       0.998    5.27
                                                    51       0.891    0.99
                                          023       11       0.866    1.14
                                                    5B       1.001    4.60
018     0.000       0.00     0.00         02A       1B       0.602    0.91
                                                    5B       0.597    0.65
                                          02B       65       0.565    0.68
                                                    2F       0.596    0.28
01A     0.091       0.12     0.18         02E       11      35.504    2.24
                                          02F       51      34.295    6.20
                                                    65      22.150   29.75
023     0.016       0.03     0.00         0FC       14       9.742    3.28
                                                    54       6.393   31.57
                                          0FD       24      10.991    3.64
                                                    64       6.117   20.95
024     0.197       0.22     0.00         0FE       2E      11.473   16.29
                                                    6E      14.945    5.10
                                          0FF       1D       9.308   26.27
                                                    5D      10.643   16.64
028     0.007       0.00     0.00         040       02       0.888    0.37
                                                    42       0.874    0.38
                                          041       16       0.918    0.24
                                                    56       0.900    0.37
029     0.000       0.00     0.01         04C       24       0.525    0.11
                                                    5A       0.505    0.11
                                          04D       2E       0.511    0.22
                                                    6E       0.471    0.00
035     0.000       0.00     0.21         600       28       7.583    0.49
                                          601       68       7.564    0.53
```

Figure 3-25. I/O Queuing Activity Interval Report (3090)

The report for the 3090 processor includes three fields that present information about the I/O processor and initiative queuing as well as information about the effective use of channel paths for individual LCUs.

In a PR/SM environment in LPAR mode, the top section of the report, which reports TOTAL SAMPLES, IOP, ACTIVITY RATE and AVG Q LENGTH, measures PR/SM system activity. The rest of the report applies to I/O activity for the MVS system being measured by RMF.

Figure 3-26 (Page 1 of 2). Fields in the I/O Queuing Activity Report (3090)

Field Heading	**Definition**
IOP	The two digit hexadecimal representation of the I/O processor (IOP). There may be more than one IOP for a 3090 processor. In a PR/SM environment, this value represents overall PR/SM system activity.
ACTIVITY RATE	The rate per second at which the IOP places I/O requests on the IOP initiative queue. There is one initiative queue for each IOP, and this value reflects the load of I/O requests on each IOP. The calculation used to compute the activity rate is: `AR = TR / SI` The activity rate (AR) equals the total number of requests on the initiative queue (TR) divided by the number of seconds in the interval (SI). In a PR/SM environment, this value represents overall PR/SM system activity.
AVG Q LNGTH	The average number of entries on the initiative queue for this IOP. Each time the IOP adds a request to the initiative queue, the current queue length is added to an accumulator. At the end of the interval, RMF divides the contents of the accumulator by the number of times an element was added to the queue. The calculation used is: `AQL = (QL-TR) / TR` The average queue length (AQL) for the interval equals the accumulated queue length for the initiative queue (QL) minus the total number of requests enqueued on the initiative queue (TR) divided by the total number of requests enqueued on the initiative queue (TR). In a PR/SM environment, this value represents overall PR/SM system activity.
LCU	The three-digit hexadecimal number that identifies the logical control unit (LCU). The set of devices associated with an LCU measurement are model-dependent. On a 3090 processor, an LCU is the logical representation of a physical control unit or a group of physical control units with one or more devices in common. Each physical control unit and each device can belong to only one LCU; they cannot be shared between LCUs. To find the LCU number, RMF must access the I/O configuration data set (IOCDS) in a 3090 processor. If RMF cannot read it, or if it has been updated so that the data might not apply to the present configuration, RMF ends the I/O queuing activity report. If no activity has occurred during the interval for all selected LCUs, the message NO ACTIVITY FOR SELECTED LCUs appears instead of the data after the headings of the report.
CONTENTION RATE	The rate per second at which the I/O processor (IOP) places delayed I/O requests on the CU-HDR for this LCU. For the 3090 processor, the IOP places an I/O request on the CU-HDR when all paths to the subchannel are busy and at least one path to the control unit is busy. For devices with only one path, or for devices where multiple paths exist and the busy condition is immediately resolved via an alternate path, the IOP does not count the condition. The calculation used is: `CR = ER / SI` The contention rate (CR) equals the total number of enqueued requests (ER) divided by the number of seconds in the interval (SI).

Figure 3-26 (Page 2 of 2). Fields in the I/O Queuing Activity Report (3090)

Field Heading	Definition
DELAY Q LNGTH	The average number of delayed requests on the control unit header (CU-HDR). Each time a request is enqueued on the CU-HDR, RMF counts the number of requests on the queue and adds that number to the accumulator. The calculation used is: `DQL = (AL-ER) / ER` The delayed queue length (DQL) equals the accumulated queue length (AL) minus the total number of enqueued requests (ER) divided by the total number of enqueued requests (ER).
% ALL CH PATH BUSY	The percentage of time during the measurement interval when all channel paths belonging to the logical control unit were busy at the same time. Unlike the corresponding field in the report for the 308x and 4381 processors, this field is not based on sampling. The data is a mathematical probability derived from multiplying the percent busy of each channel path involved. The field is blank if any vary channel path or configure channel path activity occurred during the interval. Only online channel paths that are both online to the system and connected to a device are included in the the calculation. If a channel path was online to the system but was not connected to any device or group of devices of the LCU, it will be omitted from the calculation. The calculation used is: `ACB = CHPID0 x CHPID1 x CHPID2 x CHPID3` The percentage of time during the interval when all paths in the logical control unit were busy (ACB) is calculated by multiplying the percent busy of each channel path involved (CHPID0, CHPID1, CHPID2, and CHPID3).
CONTROL UNITS	The three-digit hexadecimal identifier of each physical control unit contained in the logical control unit. There can be maximum of four physical control units associated with a 3090 logical control unit.

CHAN PATHS	The two-digit hexadecimal channel path identifier (CHPID) of the channel paths that are attached to the physical control units contained in the logical control unit. There can be up to four channel paths in a 3090 logical control unit. The channel paths that are offline or moved online or offline during the interval are indicated as follows:
	`'OFFLINE'`
	`'NOW OFFLINE'`
	`'NOW ONLINE'`
	Channel paths that are online to the system but that might or might not be connected during the interval to any device in an LCU are indicated as follows:
	`'PATH(S) OFFLINE'`
	`'PATH(S) NOW OFFLINE'`
	`'PATH(S) NOW ONLINE'`

CHPID TAKEN	The rate at which I/O requests to devices of this LCU are satisfied by each CHPID during the interval. By reviewing the rate at which each channel path of the LCU satisfies I/O requests, you can see how evenly the work requests are distributed among the available paths and how effectively those paths are arranged for the LCU. The calculation used is: `CPT = TO / SI` The rate at which I/O requests to devices of the LCU are satisfied by each CHPID (CPT) equals the total I/O operations accepted on that path (TO) divided by the seconds in the interval (SI). **Note:** If vary activity has occurred during the interval, this field is blank.
% CU BUSY	This field indicates control unit contention. Reported for each path within the logical control unit, it is the ratio of the number of requests deferred due to control unit busy and the number of attempts to service I/O requests during the measurement interval. The calculation used is: `PCB = (RCU/RA) x 100` The percentage control unit busy equals the number of requests deferred because the control unit was busy (RCU) divided by the number of I/O operations attempted (RA) multiplied by 100. Note that RA represents the sum of successful and deferred attempts to service a request on this path during the measurement interval. RMF reports a value even if the channel path changes status during the interval.

Page/Swap Data Set Activity Report

The page/swap data set activity report provides information about page/swap data set use for each individual data set. The information about the number of slots used is sampled and reported as minimum, maximum, and average values for the interval. Other information in the report gives the time when the Auxiliary Storage Manager (ASM) considered the data set to be busy, the number of start I/O requests initiated by ASM for the data set, the average service time or the average page transfer time for each I/O request, and the number of pages transferred to and from the page data set.

The information in the page data set section of the report can be used, for example, to determine whether the optimum size has been allocated for each data set. If the maximum number of slots used is consistently below the number of slots allocated, you might consider reducing the size of the data set to conserve space on the device. Use caution when reducing the size of the PLPA and common data sets, because overflow cannot occur from these data sets to the local data sets.

Use the swap data set section of the report to evaluate your allocation of swap sets. When you analyze the information on swap data sets, note that the swap set is the unit of measurement in the report. A swap set consists of twelve contiguous groups of 4096-byte slots. If the minimum and maximum swap sets consistently equal the allocated swap sets, then you might require more swap space.

The page/swap data set activity report has two major headings:

```
PAGE DATA SET USAGE
SWAP DATA SET USAGE
```

See Figure 3-27 for examples of the page/swap data set report.

Because most of the fields for these two sections are identical, they are defined together in the following section. Any differences are noted.

Duration Report: Certain fields in the page/swap data set activity duration report (DEV NUM, VOLUME SERIAL, DEVICE TYPE, SLOTS ALLOC, and SWAP SETS ALLOC) might change during a lengthy duration interval. When such a change occurs, it is not reflected in the duration report; these fields are set according to the contents of the first type 75 SMF record encountered.

P A G E / S W A P   D A T A S E T   A C T I V I T Y

PAGE 1

OS/VS2 SP2.2.3 | SYSTEM ID AQFT | RPT VERSION 3.5.1 | DATE 03/04/89 | TIME 09.00.00 | INTERVAL 29.59.997 | CYCLE 1.000 SECONDS

NUMBER OF SAMPLES = 1,800

PAGE DATA SET USAGE

PAGE SPACE TYPE	VOLUME SERIAL	DEV NUM	DEVICE TYPE	SLOTS ALLOC	---- SLOTS USED --- MIN	MAX	AVG	BAD SLOTS	% IN USE	PAGE TRANS TIME	NUMBER IO REQ	PAGES XFER'D	VIO	DATA SET NAME
COMMON	RESGPB	633	3380	3750	1665	1665	1665	0	0.00	0.000	12	7		SYS1.PLPAB
COMMON	RESGPB	633	3380	3750	17	17	17	0	0.00	0.000	0	0		SYS1.COMMONB
LOCAL	PAGE91	4B2	3380	65400	108	1243	261	0	6.50	0.005	3,601	23,349	Y	SYS1.PAGE91V
LOCAL	PAG6C1	6C1	3380	65400	116	1248	281	0	6.61	0.005	3,616	23,614	Y	SYS1.PAG6C1
LOCAL	SWP080	080	3351P	12000	3	3	3	0	0.00	0.000	0	0	N	SYS1.PAGE01
LOCAL	SWP090	090	3351P	12000	0	0	0	0	0.00	0.000	0	0	N	SYS1.PAGE03
LOCAL	PGE100	100	3351P	12000	0	0	0	0	0.00	0.000	0	0	N	SYS1.PAGE07
LOCAL	PGE110	110	3351P	12000	0	0	0	0	0.00	0.000	0	0	N	SYS1.PAGE09
LOCAL	PGE120	120	3351P	12000	0	0	0	0	0.00	0.000	0	0	N	SYS1.PAGE11
LOCAL	PGE130	130	3351P	12000	0	0	0	0	0.00	0.000	0	0	N	SYS1.PAGE13

NUMBER OF SAMPLES = 1,800

SWAP DATA SET USAGE

PAGE SPACE TYPE	VOLUME SERIAL	DEV NUM	DEVICE TYPE	SWAP SETS ALLOC	-- SWAP SETS USED-- MIN	MAX	AVG	BAD SWAP SETS	% IN USE	AVG SERV TIME	NUMBER IO REQ	DATA SET NAME
SWAP	SWP081	081	3351P	1000	0	0	0	0	0.00	0.000	0	SYS1.SWAP02
SWAP	SWP091	091	3351P	1000	0	0	0	0	0.00	0.000	0	SYS1.SWAP04
SWAP	PGE121	121	3351P	1000	0	0	0	0	0.00	0.000	0	SYS1.SWAP12
SWAP	PGE131	131	3351P	1000	0	0	0	0	0.00	0.000	0	SYS1.SWAP14
SWAP	PGE101	101	3351P	1000	0	0	0	0	0.00	0.000	0	SYS1.SWAP08
SWAP	PGE111	111	3351P	1500	0	0	0	0	0.00	0.000	0	SYS1.SWAP10

Figure 3-27. Page/Swap Data Set Activity Report

Figure 3-28 (Page 1 of 2). Fields in the Page/Swap Dataset Activity Report

Field Heading	Definition
PAGE SPACE TYPE	This field appears in both sections of the report. It indicates the page space type, which can be PLPA, COMMON, DUPLEX, LOCAL, or SWAP.
VOLUME SERIAL	This field appears in both sections of the report. It indicates the volume serial number of the volume on which the data set resides.
DEV NUM	This field appears in both sections of the report. It indicates the number of the device on which the data set resides.
DEVICE TYPE	This field appears in both sections of the report. It indicates the device type on which the data set resides.
SLOTS ALLOC	This field appears only in the **page data set section** of the report. It indicates the total number of slots each page data set contains.
SWAP SETS ALLOC	This field appears only in the **swap data set section** of the report. It indicates the total number of swap sets each swap data set contains.
----SLOTS USED --- (MIN MAX AVG)	These fields appear only in the **page data set section** of the report. They indicate the number of slots that were being used for pages. These values are sampled and reported as minimum, maximum, and average. **Note:** The actual minimum and maximum values might occur at a time when RMF is not sampling.
----SWAP SETS USED --- (MIN MAX AVG)	These fields appear only in the **swap data set section** of the report. They indicate the number of swap sets that were being used for pages. These values are sampled and reported as minimum, maximum and average. **Note:** The actual minimum and maximum values might occur at a time when RMF is not sampling.
BAD SLOTS	This field appears only in the page data set section of the report. It indicates the number of slots that encountered permanent I/O errors.
BAD SWAP SETS	This field appears only in the swap data set section of the report. It indicates the number of swap sets that encountered permanent I/O errors.
% IN USE	This field appears in both sections of the report. It indicates the percentage of time during the reporting interval when the data set was considered busy by the Auxiliary Storage Manager (ASM). At each cycle, RMF tests the busy bit in the IORB(s) for each data set. When at least one busy bit is on, RMF adds one to a counter. At the end of the interval, the percentage is computed. The calculation used is: `P = 100 x C/N` The in-use percentage equals 100 times the count of busy samples divided by the total number of samples.
AVG SERV TIME	This field appears only in the swap data set section of the report. It indicates the average number of seconds required to complete each I/O request for the data set. The calculation used is: `T = N/S x L/R` The average service time (T) equals the number of samples when the data set was in use (N) divided by the total number of samples (S) multiplied by the length of the reporting interval (L) divided by the total number of successful Start Subchannel (SSCH) instructions issued to the data set during the interval (R). **Note:** For a multiple exposure device, N is the number of samples when RMF found each exposure in use.

Figure 3-28 (Page 2 of 2). Fields in the Page/Swap Dataset Activity Report

Field Heading	Definition
PAGE TRANS TIME	This field appears only in the page data set section of the report. It indicates the average number of seconds required to complete a page transfer. The calculation used is: $T = \left(\frac{N x I}{S}\right) / R$ The average page transfer time (T) equals the number of samples when the data set was in use (N) multiplied by the length of the reporting interval (I) divided by the total number of samples (S), with the result being divided by the total number of pages transferred (R). **Note:** For a multiple exposure device, N is the number of samples when RMF found each exposure in use.
NUMBER IO REQ	This field appears in both sections of the report. It indicates the total number of I/O requests for the data set made during the interval.
PAGES XFER'D	This field appears only in the page data set section of the report. It indicates the number of pages that were transferred to or from the page data set.
DATA SET NAME	This field appears in both sections of the report. It indicates the name of the page or swap data set being monitored. A page data set name longer than 35 characters will be truncated to 35 characters in the report. A swap data set name longer than 43 characters will be truncated to 43 characters in the report. The entire data set name appears in the SMF record. **Notes:** • If a data set was dynamically introduced during the interval, its data set name is preceded by an asterisk (*). • When the operating system has detected errors in a data set that prevents its further use, the name of the data set is preceded by two asterisks (**). ASM continues to access the data set in read-only mode, and RMF reports this activity.
VIO	This field indicates whether the local paging data set accepts VIO pages. The symbols are: **Y** Indicates VIO pages are accepted **N** Indicates VIO pages are not accepted.

Paging Activity Report

The paging activity report provides information about the demands made on the system paging facilities, and the utilization of real storage, extended storage, and external page storage during the interval. It also provides swapping statistics for the interval. The paging rates monitored are:

- Non-swap page reclaim rates
- Non-swap page-in rates
- Swap page-in rates
- Total page-in rates
- Non-swap page-out rates
- Swap page-out rates
- Total page-out rates

All of the above paging data rates appear for one or more of the following:

- Pageable system areas used for non-VIO data, broken down into LPA and CSA
- Address space pages used for VIO data
- Address space pages used for non-VIO data

On a 3090 processor, the information about the use of extended storage includes the rate of page movement to extended storage and the migration rate. In addition, RMF provides information about available extended storage frames, the contention for extended storage, and the extended storage configuration.

The paging rate section provides information about page movement in real storage and the use of extended storage.

The swap information includes a description of swap placement activity. It reports data for logical swapping, physical swapping to auxiliary storage, and, for 3090 processors, physical swapping to extended storage. At the top of the page, the report provides the current OPT member. At the bottom of the report, one line indicates the average number of pages per swap-in and swap-out; a second line at the bottom of the page indicates the number of terminal output wait occurrences that SRM has detected.

For 308x and 4381 processors, the headings for the extended storage fields appear on the report, but the fields contain zeroes.

The paging rate section provides information about page movement in real storage and the use of extended storage.

The swap information includes a description of swap placement activity. It reports data for logical swapping, physical swapping to auxiliary storage. At the top of the page, the report provides the current OPT member. At the bottom of the report, one line indicates the average number of pages per swap-in and swap-out; a second line at the bottom of the page indicates the number of terminal output wait occurrences that SRM has detected.

For all processors, the frame and slot counts included in the report are:

- Total main storage page frames, categorized by LPA, CSA, private area, unused, and total
- Fixed main storage page frames, categorized by SQA, LPA, CSA, LSQA, private area, below 16 megabyte, and total
- Nucleus frames
- Total online system frames
- Total local page data set page slots
- Auxiliary slots used for VIO data set pages
- Auxiliary slots used for non-VIO private area pages
- Unused auxiliary slots
- Bad auxiliary slots (unavailable)

The data in the paging activity report can be used in a variety of ways to analyze your system's performance. For example, you can determine the non-swap, non-VIO page fault rate by adding the number of page-ins to the number of reclaims. If this is excessively high, it can be the result of an over commitment of real storage, perhaps because of a specification in the IPS of minimum multiprogramming levels that are too high for the system configuration and current workload mix.

Other problems to look for are a high pageable system area non-swap page-in rate, which could be caused by a poor pack list or a large number of fixed LPA modules. A period of high VIO slot use could be a sign that a specific job is making excessive use of VIO. And, of course, always be alert for bad slots because they can cause executing jobs to terminate abnormally.

Paging activity is formatted in four sections, each with a separate subheading:

```
MAIN STORAGE PAGING RATES PER SECOND
MAIN STORAGE FRAME COUNTS
LOCAL PAGE DATA SET SLOT COUNTS
SWAP PLACEMENT ACTIVITY
```

Figure 3-29 shows an example of the main storage paging rates section of a paging activity report.

```
                                        P A G I N G   A C T I V I T Y
                                                                                                      PAGE 1

          OS/VS2                    SYSTEM ID AQFT            DATE 03/04/89            INTERVAL 29.59.997
          SP2.2.3                   RPT VERSION 3.5.1         TIME 09.00.00            CYCLE 1.000 SECONDS
  OPT = IEAOPTFT                        MAIN STORAGE PAGING RATES PER SECOND
--------------------------------------------------------------------------------------------------------------

                                                                                        | REAL STORAGE MOVEMENT
                   PAGE RECLAIMS                 PAGE IN                     PAGE OUT   | TOTAL RATE     163.85
                   -------------   ----------------------------   ----------------------------  |-----------------------
                            %                                %                             %
                            OF                               OF                            OF   | RATE OF PAGE MOVEMENT
                   NON      TOTL            NON     TOTAL  TOTL             NON     TOTAL  TOTL | TO ES           46.15
  CATEGORY         SWAP     NON SWAP  SWAP  SWAP    RATE   SUM        SWAP  SWAP    RATE   SUM  |-----------------------

  PAGEABLE SYSTEM                                                                            | MIGRATION RATE   0.00
  AREAS (NON VIO)                                                                            |
     LPA           0.00     0               0.00    0.00   0                                 |-----------------------
                                                                                             |   AVAILABLE ESF
     CSA           0.00     0               0.00    0.00   0                0.00    0.00   0 | MIN   29465
                   ------   ----            ------  ------ ----            ------  ------ ----| MAX   42001
     SUM           0.00     0               0.00    0.00   0                0.00    0.00   0 | AVG   34282
  ADDRESS SPACES                                                                             |
     VIO           6.12     75              12.63   12.63  100              13.46   13.46  100 -----------------------
                                                                                             |  HIGH UIC  MIGR AGE
     NON VIO       2.03     25        0.00  0.00    0.00   0          0.00  0.00    0.00   0 | MIN  255    34510
                   ------   ----      ------ ------ ------ ----      ------ ------  ------ ----| MAX  255    37095
     SUM           8.15     100       0.00  12.63   12.63  100        0.00  13.46   13.46  100 | AVG  255.0  35801.6
                                                                                             | -----------------------
  TOTAL SYSTEM
     VIO           6.12     75              12.63   12.63  100              13.46   13.46  100 |  ESF CONFIGURATION
                                                                                             | INSTALLED      ONLINE
     NON VIO       2.03     25        0.00  0.00    0.00   0          0.00  0.00    0.00   0 | ---------      ------
                   ------   ----      ------ ------ ------ ----      ------ ------  ------ ----|   65536         65536
     SUM           8.15     100       0.00  12.63   12.63  100        0.00  13.46   13.46  100 |

        MAIN STORAGE FRAME COUNTS              NUMBER OF SAMPLES    LOCAL PAGE DATA SET SLOT COUNTS
  ---------------------------------------------  ----------------  ------------------------------------------------
                       MIN       MAX       AVG          181                               MIN        MAX        AVG
  TOTAL FRAMES
     LPA              1,189     1,196     1,191                   AVAILABLE SLOTS     200,992    202,564    202,253
     CSA              1,459     1,676     1,558
     PRIVATE AREA    46,820    51,924    50,264                   VIO SLOTS               233      1,805        543
     UNUSED             206     5,179     1,209
                   --------  --------  --------                   NON-VIO SLOTS             3          3          3
     TOTAL           64,339    64,339    64,339
                                                                  BAD SLOTS                 0          0          0
  FIXED FRAMES                                                                      ---------  ---------  ---------
     SQA              1,234     1,278     1,263                   TOTAL SLOTS         202,800    202,800    202,800
     LPA                 51        53        51
     CSA                158       162       159
     LSQA             8,190     9,303     8,850
     PRIVATE AREA       717       872       779
     BELOW 16 MEG       895     1,089       974
                   --------  --------  --------
     TOTAL FIXED     10,364    11,652    11,103
  NUCLEUS FRAMES AT END OF INTERVAL       1,020
  TOTAL ONLINE SYSTEM FRAMES AT END OF INTERVAL      65,359
```

Figure 3-29. Main Storage Paging Rates and Page Movement Data from a Paging Activity Report

Fields in the **MAIN STORAGE PAGING RATES PER SECOND** section of the paging activity report.

Figure 3-30 (Page 1 of 3). Paging Activity Report Fields (Main storage paging rates)

Field Heading	Definition
CATEGORY	The identification of the component parts of paging rates identifying these basic components: • Pageable system area, non-VIO data • Address space, VIO data • Address space, non-VIO data
PAGEABLE SYSTEM AREAS	Areas of main storage that are not associated with a single address space. This major category consists of the following subcategories: • LPA - all values are reported except for swaps. • CSA - all values are reported except for swaps. • SUM - SUM of LPA and CSA.
ADDRESS SPACES	Areas of main storage that are associated with individual address spaces. This major category consists of the following subcategories: • VIO All values are reported except for swaps. • NON VIO All values are reported. • SUM Sum of address space VIO and non-VIO.
TOTAL SYSTEM	A major category consisting of the sum of system pageable areas and address space values and with the following subcategories: • VIO Consists only of address space VIO values. • NON VIO Sum of system pageable areas non-VIO and memory non-VIO values. • SUM Sum of system pageable areas sum and address space sum. (The computer system grand total for paging rates.)
PAGE RECLAIMS	NON SWAP - The per-second rate of pages reclaimed exclusive of address space swap-ins. The range of values is 0.00 to 999,999. % OF TOTL SUM - The percentage of the total reclaim rate for non-swap pages. The range of values is 0 to 100 percent.
PAGE IN	SWAP - The per-second rate of pages read into real storage as a result of address space swap-ins. The range of values is 0.00 to 999,999. NON SWAP - The per-second rate of pages read into real storage from auxiliary storage exclusive of address space swap-ins. Non-VIO paging occurs as a result of a page fault, PGLOAD, or PGFIX. When there are concurrent requests for the same page, the first generates a page-in; all the rest are considered reclaims. VIO paging occurs as a result of a page fault or PGLOAD on a VIO window (logical GETs). VIO pages that are swapped in are not included. The range of values is 0.00 to 999,999 pages per second. TOTAL RATE - The per-second rate of the total system pages read into main storage. The rate is the sum of the non-swap page-in rate and the swap page-in rate. The range of values is from 0.00 to 999,999. % OF TOTL SUM - The percentage of the total page-in rate for each part of the total. The total here is reflected in the subcategory SUM under TOTAL SYSTEM. The range of values is 0 to 100 percent.

Figure 3-30 (Page 2 of 3). Paging Activity Report Fields (Main storage paging rates)	
Field Heading	**Definition**
PAGE OUT	SWAP - The per-second rate of pages written to auxiliary storage as a result of address space swap outs. The range of values is 0.00 to 999,999. NON SWAP - The per-second rate of pages written to auxiliary storage (forced out) independent of address space swap outs. Non-VIO paging results from a PGOUT (including page stealing and other RSM-generated page-outs). VIO paging results from a PGOUT (including stealing and other RSM-generated page-outs) on a VIO window page (logical PUTs). Included also are the pages trimmed at swap out for logical swap out and from address spaces protected by real storage isolation by means of the IPS minimum working set size specification. VIO pages transferred as a result of a swap-out are not included. The range of values is 0.00 to 999,999. TOTAL RATE - The per-second rate of total system pages written to auxiliary storage. The rate is the sum of the non-swap page-out rate and the swap page-out rate. The range of values is from 0.00 to 999,999. % OF TOTL SUM - The percentage of the total page-out rate for each part of the total. The total here is reflected in the subcategory SUM under TOTAL SYSTEM. The range of values is 0 to 100 percent.
REAL STORAGE MOVEMENT	TOTAL RATE - The per second rate of total page movement between storage locations above 16 megabytes and storage locations below 16 megabytes or vice versa. The rate is calculated by dividing the total number of pages moved by the length of the interval. The range of values is 0.00 to 999,999.
RATE OF PAGE MOVEMENT TO EXTENDED STORAGE	The total rate per second of pages sent to extended storage because of both paging and swapping requests. This field indicates how much use your system makes of extended storage. When RMF is running on a **308x or a 4381 processor**, this field contains zeroes.
MIGRATION RATE	The total rate per second of pages migrated from extended storage to auxiliary storage. The migration rate identifies, overall, how effectively you are using extended storage. A low migration rate means that pages that are paged out to extended storage are referenced quickly enough to avoid migration to an external I/O paging device. A low migration rate means that the system is avoiding a high paging overhead. A high migration rate, in contrast, might mean contention for limited extended storage or that the extended storage criteria table does not match your workload. When RMF is running on a **308x or a 4381 processor**, this field contains zeroes.
AVAILABLE ESF (MIN MAX AVG)	The number of extended storage frames actually available out of the total extended storage installed for the system. This value indicates the amount of unused extended storage and will tend to vary over the measurement interval; therefore, RMF provides minimum, maximum, and average values. The AVAILABLE ESF field contains a one to four character value expressed as follows: 0 - 9999 from 0 to 9,999 frames 10K - 999K from 10,000 to 999,000 frames 1M - 999M from 1,000,000 to 999,000,000 frames When RMF is running on a **308x or a 4381 processor**, this field contains zeroes.
HIGH UIC (MIN MAX AVG)	The system high unreferenced interval count (H UIC). The HIGH UIC can indicate whether there is contention for real storage frames. When the HIGH UIC is relatively low, contention for real storage is high. Although total paging rates might vary with the type and level of workload, the unreferenced interval count is the best indicator of actual storage contention. Because this value varies over the measurement interval, RMF provides minimum, maximum and average HIGH UIC values.

Figure 3-30 (Page 3 of 3). Paging Activity Report Fields (Main storage paging rates)	
Field Heading	**Definition**
MIGR AGE	The length of time a page remains unreferenced in extended storage before it migrates to auxiliary storage. A high migration rate implies a low migration age, and both measurements can indicate an over-use of extended storage. On the other hand, a low migration rate and a high migration age could mean under-use of extended storage. By comparing the migration age to the criteria age of each category on the criteria table, you can understand which types of pages are being sent to extended storage and which are not. Because this value varies over the measurement interval, RMF provides minimum, maximum and average migration age values. When RMF is running on a **308x or a 4381 processor**, this field contains zeroes.
ESF CONFIGURATION	INSTALLED EXTENDED STORAGE FRAMES - the total amount of installed extended storage. ONLINE EXTENDED STORAGE FRAMES - the amount of installed extended storage that is actually online to the system. When RMF is running on a **308x or a 4381 processor**, this field contains zeroes.

Fields under the **MAIN STORAGE FRAME COUNTS** section of the paging activity report

Figure 3-31. Paging Activity Report Fields (Main Storage Frame Counts Section)	
Field Heading	**Definition**
TOTAL FRAMES	This data is sampled and presented as MIN (minimum), MAX (maximum), and AVG (average) values. This particular category of main storage is further broken down into LPA, CSA, PRIVATE AREA, UNUSED, and TOTAL. The TOTAL value is not a summation of the previous four categories at the end of the interval, but is derived by adding the four categories during each sample and later reporting the MIN sum, MAX sum, and AVG sum for the set of samples. The total frame count includes frames both above and below 16 megabytes. **Note:** The actual maximum or minimum value of a field might occur at a time when RMF is not sampling.
FIXED FRAMES	This data is sampled and presented as MIN, MAX, and AVG values. This category is further broken down into SQA, LPA, CSA, LSQA, PRIVATE AREA, BELOW 16 MEG, and TOTAL FIXED. The TOTAL FIXED value is not a summation of the previous six categories at the end of the interval but is derived by adding the categories during each sample and reporting the MIN sum, MAX sum, and AVG sum for the set of samples. This section of the paging activity report also provides end-of-interval values for nucleus frames and total online system frames. The nucleus frame count includes both the DAT-ON nucleus and the DAT-OFF nucleus.
NUCLEUS FRAMES AT END OF INTERVAL	The number of frames at the end of the interval that includes the read/write nucleus, the read only nucleus, and the DAT-OFF nucleus.
NUMBER OF SAMPLES	The number of samples taken in this interval. SMF record type 71 (paging activity) contains data about skipped samples obtained as a result of a non-zero return code from IARXCNTF (the RSM fixed-frame count interface). If the number of valid samples obtained via IARXCNTF does not equal the number of samples in SMF71SAM, this column contains the value determined by IARXCNTF highlighted by a series of asterisks as follows: `*** number of valid samples obtained via IARXCNTF ***` See *MVS/XA SPL: System Management Facilities (SMF)*.

Fields in the **LOCAL PAGE DATA SET SLOT COUNTS** section of the paging activity report.

This data is sampled and presented as MIN, MAX, and AVG values as described previously. This category of storage is further broken down into:

Figure 3-32. Paging Activity Report Fields (Local Page Data Set Slot Counts Section)

Field Heading	Definition
AVAILABLE SLOTS	The number of page slots that do not contain any data pages and that are available for use.
VIO SLOTS	The number of local page data set slots that contain pages for VIO data sets.
NON-VIO SLOTS	The number of local page data set slots that contain pages that belong to address-space virtual storage.
BAD SLOTS	The number of local page data set slots that do not contain any data pages and are unavailable for use because of permanent I/O errors.
TOTAL SLOTS	This value is not a summation of the previous four categories at the end of the interval; it is derived as the TOTAL FIXED frame value is derived.

SWAP PLACEMENT ACTIVITY

The swap placement activity section of the paging activity report provides a full analysis of swap placement. It includes information about extended storage for 3090 processors.

To help you analyze the data, the current OPT member appears after the report headings. The OPT member contains the current extended storage criteria table used to tune the use of extended storage.

This section of the report identifies by reason the total swap occurrences (logical swaps and physical swaps) and gives an analysis of the resulting swap transitions. The initial swap decisions that involve extended storage on a 3090 processor are:

1. **Logical swapping in real storage** - Of the total number of swap candidates, some may be logically swapped in real storage. At some later time, some of these may be physically swapped to extended storage or physically swapped to auxiliary storage. Of those physically swapped to extended storage, some may be migrated to auxiliary storage at a later time.
2. **Physical swapping to extended storage** - Of the total number of logical swap candidates, some may be physically swapped directly to extended storage. Of these, some may be migrated to auxiliary storage at a later time.
3. **Physical swap to auxiliary storage** - Of the total number of logical swap candidates, some may be placed directly on auxiliary storage.

Figure 3-33 shows an example of the swap placement activity section of the paging activity report.

```
                                          P A G I N G   A C T I V I T Y
                                                                                                  PAGE 1

           OS/VS2                   SYSTEM ID AQFT          DATE 03/04/89          INTERVAL 29.59.997
           SP2.2.3                  RPT VERSION 3.5.1       TIME 09.00.00          CYCLE 1.000 SECONDS

OPT = IEAOPTFT                               S W A P   P L A C E M E N T   A C T I V I T Y
                             *------ AUX STORAGE -------* *---LOGICAL SWAP--*  *---------------- EXTENDED STORAGE ----------------*
                             AUX               AUX STOR                                                MIGRATED               LOG SWAP
                             STOR    AUX STOR   VIA                 LOG SWAP   EXT STOR   EXT STOR      FROM     EXT STOR   /EXT STOR
                     TOTAL   TOTAL   DIRECT    TRANSITION LOG SWAP  EFFECTIVE  DIRECT      TOTAL      EXT STOR   EFFECTIVE  EFFECTIVE
TERMINAL      CT    18,818      0         0          0    18,810    18,442        8         376           0        376      18,818
INPUT/OUTPUT  RT     10.45   0.00      0.00       0.00     10.45     10.25     0.00        0.21        0.00       0.21       10.45
WAIT           %     94.2%   0.0%      0.0%       0.0%    100.0%     98.0%     0.0%        2.0%        0.0%     100.0%      100.0%
LONG          CT       453      0         0          0       453       453        0           0           0          0         453
WAIT          RT      0.25   0.00      0.00       0.00      0.25      0.25     0.00        0.00        0.00       0.00        0.25
               %      2.3%   0.0%      0.0%       0.0%    100.0%    100.0%     0.0%        0.0%        0.0%       0.0%      100.0%
DETECTED      CT       668      0         0          0       668       668        0           0           0          0         668
WAIT          RT      0.37   0.00      0.00       0.00      0.37      0.37     0.00        0.00        0.00       0.00        0.37
               %      3.3%   0.0%      0.0%       0.0%    100.0%    100.0%     0.0%        0.0%        0.0%       0.0%      100.0%
UNILATERAL    CT        38      0         0          0        38        37        0           1           0          1          38
              RT      0.02   0.00      0.00       0.00      0.02      0.02     0.00        0.00        0.00       0.00        0.02
               %      0.2%   0.0%      0.0%       0.0%    100.0%     97.4%     0.0%        2.6%        0.0%     100.0%      100.0%
EXCHANGE ON   CT         7      0         0          0         7         7        0           0           0          0           7
RECOMMENDA-   RT      0.00   0.00      0.00       0.00      0.00      0.00     0.00        0.00        0.00       0.00        0.00
TION VALUE     %      0.0%   0.0%      0.0%       0.0%    100.0%    100.0%     0.0%        0.0%        0.0%       0.0%      100.0%
TOTAL         CT    19,984      0         0          0    19,976    19,607        8         377           0        377      19,984
              RT     11.10   0.00      0.00       0.00     11.10     10.89     0.00        0.21        0.00       0.21       11.10
               %    100.0%   0.0%      0.0%       0.0%    100.0%     98.2%     0.0%        1.9%        0.0%     100.0%      100.0%




AUXILIARY STORAGE - AVERAGE PAGES PER SWAP OUT -         0       AVERAGE PAGES PER SWAP IN -         0
OCCURRENCES OF TERMINAL OUTPUT WAIT -     1,302
```

Figure 3-33. Swap Placement Activity of Paging Activity Report

The entries in the leftmost column of the swap placement activity section identify the swap reason. For each swap reason, the report provides the count (CT), rate (RT), and percentage (%). The swap reasons are:

Figure 3-34 (Page 1 of 2). Common Swap Reasons in the Paging Activity Report

Swap Reason	**Definition**
TERMINAL INPUT/OUTPUT WAIT	Number of swap-outs that occur because a terminal was waiting for input or was waiting for output buffers, but the address space was not a candidate for logical swapping. Although TERMINAL INPUT/OUTPUT WAIT is identified as a single reason, SRM distinguishes INPUT/OUTPUT as two separate swap reasons, each having its unique set of counters. To derive the fields for TERMINAL INPUT/OUTPUT WAIT, RMF always uses the sum of both counters in its calculation.
LONG WAIT	Number of swap-outs that occur because of user-requested long waits. Swap-outs of address spaces that were swapped in, executed, and entered a wait state before SRM received control are also counted.
DETECTED WAIT	Number of swap-outs that occur because of user address spaces that have not issued a WAIT,LONG=YES macro but have gone without executing for more than the SRM-specified wait threshold or 2 seconds, whichever is greater. The SRM determines the wait threshold by dividing 8 seconds by the CPU-related SRM second/real time seconds.
UNILATERAL	Number of swap-outs that occur because the multiprogramming level for a domain became higher than the target multiprogramming level specified for that domain.

Figure 3-34 (Page 2 of 2). Common Swap Reasons in the Paging Activity Report	
Swap Reason	**Definition**
EXCHANGE ON RECOMMENDATION VALUE	Number of swap-outs that occur because one user in a domain must be swapped out to allow another user (with a higher workload level) in the same domain to be swapped in.

The following swaps rarely occur:

Figure 3-35. Uncommon Swap Reasons in the Paging Activity Report	
ENQUEUE EXCHANGE	Number of swap-outs that occur to make room for users enqueued upon a system resource that is required by other users.
REQUESTED	Number of swap-outs that occur because the program properties table (PPT) indicated that a job or step is non-swappable. The address space is swapped out, and its LSQA is swapped into non-reconfigurable storage. This field also includes the number of swap outs caused by a CONFIG STOR,OFFLINE command.
AUXILIARY STORAGE SHORTAGE	Number of swap-outs that occur because 70% of all local page slots are allocated. The swappable user address spaces acquiring auxiliary storage page slots at the greatest rate are swapped out.
REAL PAGEABLE STORAGE SHORTAGE	Number of swap-outs that occur because of a shortage of real pageable frames. The swappable users that have acquired the most fixed pages are swapped out.
TRANSITION TO NONSWAPPABLE	Number of swap-outs that occur because an address space is made non-swappable. Bits are set when the address space is made non-swappable to ensure that it will use the preferred area and not the reconfigurable area after it is made non-swappable. If no swaps occurred during the measurement interval for a particular swap reason, the report omits that swap reason and its associated counts, rates and percentages. For example, few, if any, swaps usually occur for the following swap reasons: • Requested swap • Auxiliary storage shortage • Real pageable storage shortage • Transition to nonswappable state

The following categories appear in the swap placement activity report:

Figure 3-36. Fields in the Paging Activity Report (Swap Sequence Counts)

Field Heading	Definition
TOTAL	The swap requests for this reason in relation to the total number of swap requests. For each reason, it describes the total number of swap requests, the rate per second, and the percentage of all swaps that occurred for that reason.
AUX STOR TOTAL	The total number of swap requests directly or indirectly placed on auxiliary storage.
AUX STOR DIRECT	The swap requests for this reason in relation to the total number of swap requests originally directed to auxiliary storage. Percent data is a percentage of the AUX STOR TOTAL.
AUX STOR VIA TRANSITION	The swap requests for this reason in relation to the total number of swap requests eventually placed on auxiliary storage. These were either originally logically swapped from real storage to auxiliary storage or directed from real storage to extended storage. Percent data is a percentage of the AUX STOR TOTAL.
LOG SWAP	The logical swap candidates that were logically swapped.
LOG SWAP EFFECTIVE	The logical swap candidates that were logically swapped and never physically swapped. Percent data is a percentage of LOG SWAP.
EXT STOR DIRECT	The logical swap candidates originally directed to extended storage.
EXT STOR TOTAL	The total of all swap requests directed to extended storage. These include those swaps originally directed to extended storage and those logically swapped that were physically swapped and directed to extended storage.
MIGRATED FROM EXTENDED STORAGE	Out of the total number of swap requests placed on extended storage, those eventually moved to auxiliary storage. The percent field is a sum of the percentage of EXT STOR TOTAL and MIGRATED FROM EXTENDED STORAGE.
EXT STOR EFFECTIVE	Those swap requests that are originally directed to extended storage and not migrated to auxiliary storage. The percent field is a percentage of EXT STOR TOTAL.
LOG SWAP/EXT STOR EFFECTIVE	Those swap requests that either remained logically swapped in real storage or, after being directed to extended storage, remained there and were not migrated. This value indicates that I/O to auxiliary storage was avoided when the swap request was satisfied. Logically, this value is the sum of LOG SWAP EFFECTIVE plus EXT STOR EFFECTIVE.
AVERAGE PAGES PER SWAP OUT	The average number of pages swapped out for each storage swap-out. The range of values is 0 to 4,096 pages or 0 to 16,777,216 bytes.
AVERAGE PAGES PER SWAP IN	The average number of pages swapped in for each storage swap-in.
OCCURRENCES OF TERMINAL OUTPUT WAIT	The number of terminal output wait occurrences detected by SRM. Column headings (from left to right in the report) describe how SRM handles swap requests and how effective SRM is in the way it satisfies these requests. If extended storage is not available or not used, the rightmost five column headings still appear, but data fields for count, rate and percentage are blank. **Note:** Unless otherwise noted, percent data is always a percentage of the corresponding TOTAL.

Virtual Storage Activity Report

The virtual storage activity report provides information about the use of virtual storage that can help your installation to manage its use of virtual storage.

Information on virtual storage use is particularly helpful in the process of long-term measurement. It helps you, for example, to understand your current use of virtual storage, see the relationship between increased use of your system and increased demands on virtual storage, and predict future constraints before they occur. This ability to predict a future constraint is particularly useful for the virtual storage resource because actions that can relieve a virtual storage constraint generally require significant time to plan and implement. And, of course, the report can also help you to determine the effect of any actions, such as moving a large application above the 16-megabyte line or installing products that take advantage of extended addressing.

The information in the report can also help you to identify any expansion of SQA into CSA and to set appropriate size values for CSA and SQA at IPL time. You can also use the report to verify the cost (in increased PLPA intermodule space) of any pack lists your installation uses to reduce PLPA paging.

Most of the information you need to begin to manage virtual storage appears in the common storage summary report. When this report indicates a problem, you can request the common storage detail report or the private area summary or detail report to obtain additional information. The contents of each report are described later in this section.

All of the reports, however, define virtual storage space as either free storage or allocated storage. *Free storage* is any block of at least 4K (4096 bytes) that contains no storage obtained via the GETMAIN macro instruction. *Allocated storage* is any block of at least 4K that contains any storage obtained via the GETMAIN macro instruction. Thus, for the purposes of the report, free storage within a 4K block assigned to a subpool is considered to be allocated storage. Both free storage and allocated storage are reported as multiples of 4K on the reports.

All size data values are reported in bytes. The size is followed by a K (indicating the number of kilobytes the value represents) unless the size is greater that 9999K. When the size is greater than 9999K, the size is followed by an M (indicating the number of megabytes the value represents). Because peak values are especially important when analyzing virtual storage use, the minimum, maximum and average values are reported whenever useful, and the minimum and maximum values are time-stamped.

To minimize its overhead, RMF does not sample virtual storage data at every cycle. It takes one sample of virtual storage data for every ten RMF cycles. For example, if the RMF cycle is one second, RMF samples virtual storage data every ten seconds. In this case, RMF provides time stamps (accurate to within a ten-second range) for each minimum and maximum value on the report. The time stamp shows the time when RMF first observed the minimum or maximum value in the sample. Note that your installation can change this 1:10 ratio; see "Changing the Sample to Cycles Ratio" on page 3-91.

The Common Storage Summary Report

The common storage summary report enables you to measure the use of virtual storage with minimal overhead. It contains the information you need to understand your current use of virtual storage. If you archive the data, you can use differences over time to predict the need for actions to relieve a problem or constraint before it becomes critical. It also helps you to verify the size values set for CSA and SQA at IPL time and determine if you are using common storage effectively.

Figure 3-37 shows the format of the common storage summary report. Because RMF does not sample virtual storage data at every cycle, the value reported for NUMBER OF SAMPLES is less than the number of cycles. Each data field is defined in Figure 3-38 on page 3-85.

```
                                         V I R T U A L    S T O R A G E    A C T I V I T Y
                                                                                                         PAGE 1

          OS/VS2                    SYSTEM ID AQFT            DATE 03/04/89             INTERVAL 29.59.997
          SP2.2.3                   RPT VERSION 3.5.1         TIME 09.00.00             CYCLE 1.000 SECONDS

                                                    COMMON STORAGE SUMMARY
NUMBER OF SAMPLES      180
      STATIC STORAGE MAP                                            ALLOCATED CSA/SQA
  AREA        ADDRESS   SIZE           ------------ BELOW 16M -------------  ------- EXTENDED (ABOVE 16M) --------
EPVT          2A00000  2006M              MIN             MAX            AVG    MIN              MAX            AVG
ECSA          2049000  9948K         SQA   660K 09.00.00   708K 09.28.40  690K  3860K 09.00.00  3968K 09.26.00  3930K
EMLPA         203E000    44K         CSA  1932K 09.00.00  1976K 09.27.30 1957K  4284K 09.01.11  5136K 09.22.21  4679K
EFLPA               0     0K
EPLPA         1C32000  4144K         ALLOCATED CSA BY KEY
ESQA          1362000  9024K         0     396K 09.00.00   400K 09.03.10  398K  1080K 09.01.11  1096K 09.23.20  1091K
ENUC          1000000  3461K         1     240K 09.00.00   268K 09.27.30  252K  1956K 09.00.00  1956K 09.00.00  1956K
----- 16 MEG BOUNDARY ------         2       0K 09.00.00     0K              0K     0K 09.00.00     0K              0K
NUCLEUS        F68000   605K         3       0K 09.00.00     0K              0K     0K 09.00.00     0K              0K
SQA            EA8000   768K         4       0K 09.00.00     0K              0K     0K 09.00.00     0K              0K
PLPA           C56000  2376K         5       8K 09.00.00    12K 09.05.20    8K    88K 09.00.00    88K 09.00.00    88K
FLPA                0     0K         6    1060K 09.00.00  1076K 09.20.11 1070K   720K 09.00.10   848K 09.26.00   798K
MLPA           C51000    20K         7     228K 09.00.00   228K 09.00.00  228K   440K 09.01.11  1164K 09.07.20   745K
CSA            800000  4420K         8-F     0K 09.00.00     0K              0K     0K 09.00.00     0K              0K
PRIVATE          1000  8188K
PSA                 0     4K         SQA EXPANSION INTO CSA
                                             0K 09.00.00     0K              0K     0K 09.00.00     0K              0K
PLPA INTERMODULE SPACE -    6K IN PLPA AND    11K IN EPLPA
PLPA SPACE REDUNDANT WITH MLPA/FLPA -    0K IN PLPA AND    0K IN EPLPA

                    ---------- BELOW 16M ----------------  -------------- ABOVE 16M -----------
                      MIN             MAX            AVG    MIN              MAX            AVG
CSA
 FREE PAGES (BYTES)  2444K 09.27.30  2488K 09.00.00  2462K  4812K 09.22.21  5664K 09.01.11  5268K
 LARGEST FREE BLOCK  2444K 09.27.30  2484K 09.00.00  2460K  4812K 09.22.21  5184K 09.01.00  4914K
 ALLOCATED AREA SIZE 1936K 09.00.00  1976K 09.27.30  1959K  4764K 09.01.00  5136K 09.22.21  5033K
SQA
 FREE PAGES (BYTES)    60K 09.28.40   108K 09.00.00    77K  5056K 09.26.00  5164K 09.00.00  5093K
 LARGEST FREE BLOCK    60K 09.28.40    80K 09.00.00    72K  5056K 09.26.00  5152K 09.00.00  5091K
 ALLOCATED AREA SIZE  688K 09.00.00   708K 09.28.40   695K  9024K 09.00.00  9024K 09.00.00  9024K
MAXIMUM POSSIBLE USER REGION -  8040K BELOW AND 1997M ABOVE
```

Figure 3-37. The Common Storage Summary Report

Figure 3-38. Virtual Storage Activity Report, Common Storage Summary

Field Heading	Definition
STATIC STORAGE MAP	Identifies the major storage areas both above and below the 16-megabyte line. It includes the name of each area, the address of its lower boundary, and its size, reported in bytes.
ALLOCATED CSA/SQA	Presents MIN, MAX, and AVG values for allocated CSA and SQA, both below and above the 16-megabyte line. RMF calculates each size by adding the number of bytes assigned to each SQA or CSA subpool. The report also breaks down allocated CSA by key.
SQA EXPANSION INTO CSA	Reports the MIN, MAX, and AVG size of any expansion of SQA into CSA.
PLPA INTERMODULE SPACE	Reports the amount of unused space between the modules in both the PLPA and the EPLPA (the extended PLPA). If your installation uses a pack list (in the IEAPAK00 parmlib member), the values reported can help you to determine the cost of your packing algorithm in relation to its benefit, a reduction in LPA paging rates, as shown in the paging report.
PLPA SPACE REDUNDANT WITH MLPA/FLPA	Reports, for PLPA, the amount of space occupied by modules that also exist in (E)MLPA and/or (E)FLPA. For EPLPA, reports the amount of space occupied by modules that also exist in (E)MLPA or (E)FLPA.
FREE PAGES (BYTES)	Presents MIN, MAX, and AVG values, in bytes, for the amount of free storage.
LARGEST FREE BLOCK	Presents MIN, MAX, and AVG values, in bytes, for the size of the largest free block. The size of the largest free block, when compared to the total amount of free storage, is a measure of fragmentation within the common storage area. For example, when the size of the largest free block is close to the size of free storage, there is little fragmentation. The size of the largest free block is also the size of the largest GETMAIN that the system can currently satisfy within CSA or SQA.
ALLOCATED AREA SIZE	Presents the MIN, MAX, and AVG values, in bytes, for the size of the allocated area. RMF calculates this value as the difference between the highest and lowest address occupied by allocated storage, which includes all free blocks that lie between allocated blocks. Because free blocks between allocated blocks cause an increase in the virtual address range needed to hold the allocated blocks, consider this value when determining the size of CSA (and ECSA) and SQA (and ESQA). Significant fragmentation causes this number to be much larger than the amount of storage actually used.
MAXIMUM POSSIBLE USER REGION	This field reports the largest size specified on the REGION= JCL parameter that this system can satisfy (assuming a minimal number of DD statements). RMF determines this value, reported for both below and above the 16-megabyte line, by examining its own private area; it calculates the size by finding the difference between the bottom of its allocated area at RMF initialization and the start of the user region. There is no guarantee that a particular job can obtain a region of the reported size. For a job that runs under an initiator (as opposed to a job that runs as a started task), the storage the initiator obtains reduces the size of the region that the job can obtain. The number of DD statements in the JCL can also reduce the size of the region. The reported region size, however, can indicate whether a job with a known region requirement is likely to obtain the region it requires under the system conditions reflected in the report.

Common Storage Detail Report

The optional common storage detail report contains additional information about the use of CSA and SQA below the 16-megabyte line.

Figure 3-39 shows a sample printed report of common storage (detail). Each field is defined in Figure 3-40.

```
                                                    COMMON STORAGE DETAIL
ALLOCATED CSA BY SUBPOOL BY KEY (BELOW 16 MEG)                                ALLOCATED SQA BY SUBPOOL (BELOW 16M)
           SUBPOOL 227     SUBPOOL 228      SUBPOOL 231      SUBPOOL 241
                                                                              SUBPOOL   MIN              MAX              AVG
------------------------- MINIMUM ----------------------------------
 0        16K 09.00.00    4K 09.00.00     92K 09.00.00    284K 09.00.00       226        76K 09.00.00     92K 09.14.00     85K
 1         0K 09.00.00   12K 09.00.00     40K 09.00.00    188K 09.00.00       239       156K 09.00.00    160K 09.20.20    157K
 2                                                                            245       424K 09.01.11    460K 09.21.41    448K
 3
 4
 5                        4K 09.00.00                       4K 09.00.00
 6        84K 09.00.00                   964K 09.00.00     12K 09.00.00
 7                        8K 09.00.00     60K 09.00.00    160K 09.00.00
8-F
ALL      100K 09.00.00   28K 09.00.00   1156K 09.00.00    648K 09.00.00
------------------------- MAXIMUM ----------------------------------
 0        16K 09.00.00    4K 09.00.00     92K 09.00.00    288K 09.03.10
 1         4K 09.03.10   12K 09.00.00     60K 09.01.40    204K 09.22.01
 2
 3
 4
 5                        4K 09.00.00                       8K 09.05.20
 6        88K 09.06.50                   976K 09.20.11     12K 09.00.00
 7                        8K 09.00.00     60K 09.00.00    160K 09.00.00
8-F
ALL      108K 09.14.00   28K 09.00.00   1180K 09.27.30    668K 09.22.01
------------------------- AVERAGE ----------------------------------
 0        16K             4K              92K             286K
 1         0K            12K              42K             197K
 2
 3
 4
 5                        4K                                4K
 6        87K                            971K              12K
 7                        8K              60K             160K
8-F
ALL      103K            28K            1166K             659K
```

Figure 3-39. Common Storage Detail Report

Figure 3-40. Fields in the Common Storage Detail Report

Field Heading	Definition
ALLOCATED CSA BY SUBPOOL BY KEY	A matrix that presents MINIMUM, MAXIMUM, and AVERAGE use of CSA by subpools 227, 228, 231, and 241, broken down by storage key and summed for ALL keys.
ALLOCATED SQA BY SUBPOOL	Shows MIN, MAX, and AVG values for subpools 226, 239, and 245.

Private Area Summary Report

The optional private area summary report presents information about how a specific address space is using its private virtual storage. RMF uses the job name you specify when you request the report to identify the address space. In choosing a job, note that gathering data for a specific address space requires additional RMF overhead. Also, RMF can gather private area data only when a job is active at the beginning of the interval, and various conditions can limit RMF's ability to report complete private area data. These conditions, and the actions RMF takes, are described later in this section under "Partial Private Area Data."

In general, RMF can gather meaningful data only for jobs that run for a relatively long period of time. Note also that you cannot monitor the master scheduler address space.

Figure 3-41 shows a sample printed report for the private area (summary). The fields are defined in Figure 3-42, Figure 3-43, and Figure 3-44.

```
                                   V I R T U A L   S T O R A G E   A C T I V I T Y
                                                                                               PAGE 1

          OS/VS2                   SYSTEM ID AQFT          DATE 03/04/89            INTERVAL 29.59.997
          SP2.2.3                  RPT VERSION 3.5.1       TIME 09.00.00            CYCLE 1.000 SECONDS

                                               PRIVATE AREA SUMMARY
        JOB NAME -       CATALOG                                   REGION REQUESTED                0K
        STEP NAME -                                                REGION ASSIGNED (BELOW 16M)  8172K
        PROGRAM NAME -   IGG0CLX0                                  REGION ASSIGNED (ABOVE 16M)  2006M
        NUMBER OF SAMPLES -   180

                                             PRIVATE STORAGE MAP
                      BELOW 16M                                     EXTENDED (ABOVE 16M)
             7FFFFF ________________________                     ________________________ 7FFFFFFF
                    | LSQA/SWA             |                     | LSQA/SWA             |
                    | 229/230         1260K|    BOTTOM OF        | 229/230         9348K|
            6C5000  |_____09.20.20_________| ALLOCATED AREA      |_____09.04.50_________| 7F6DF000
                    | UNUSED             0K|                     | UNUSED             0K|
            800000  |______________________|  GETMAIN LIMIT      |______________________| 7FFFFFFF
                    | UNUSED          6268K|                     | UNUSED          1992M|
             A6000  |_____09.08.01_________|    TOP OF           |_____09.28.00_________|  2EFD000
                    | USER                 | ALLOCATED AREA      | USER                 |
                    | REGION           644K|                     | REGION               |
              5000  |______________________|                     |                 5108K|
                    | SYSTEM REGION     16K|                     |                      |
              1000  ------------------------                     ------------------------  2A00000

                    ---------- BELOW 16M ----------------     --------------- ABOVE 16M -----------
                      MIN             MAX             AVG    MIN             MAX             AVG
  LSQA/SWA/229/230
   FREE PAGES (BYTES)   0K 09.00.01    0K              0K  1992M 09.28.20  1993M 09.00.11  1992M
   LARGEST FREE BLOCK   0K 09.00.01    0K              0K  1992M 09.28.20  1992M 09.00.01  1992M
   PAGES ALLOCATED
     (IN BYTES)     1216K 09.00.01  1256K 09.08.01  1227K  8976K 09.29.11  9324K 09.13.50  9293K
  0USER REGION
   FREE PAGES (BYTES) 6276K 09.08.01  6340K 09.00.01  6326K     0K 09.00.01     0K             0K
   LARGEST FREE BLOCK
     IN GETMAIN LIMIT 6272K 09.08.01  6332K 09.00.01  6320K     0K 09.00.01     0K             0K
   PAGES ALLOCATED
     (IN BYTES)       616K 09.00.01   640K 09.08.01   618K  4524K 09.00.01  5028K 09.28.00  4850K
```

Figure 3-41. Private Area Summary Report

Figure 3-42. Subheadings in the Private Area Summary Report

Field Heading	Definition
JOBNAME and PROGRAM NAME	Identify the job analyzed in the report. RMF takes the program name from the JSCBPGMN (the same field SMF uses for job step records, types 4 and 30).
NUMBER OF SAMPLES	Shows the number of samples RMF used to generate the data in the report. If the job was swapped out at a time when RMF tried to sample virtual storage data, this number will be less than the number of samples reported for the common storage summary report.
REGION REQUESTED	Shows either (1) the values specified for the REGION= parameter on the JOB or EXEC JCL statement for the job step or (2) the system default used for the job step.
REGION ASSIGNED	For virtual storage below and above the 16-megabyte line, shows the region assigned to the job by installation control (the IEFUSI installation exit or IEALIMIT). This value limits the amount of storage that a job can obtain by issuing a variable length GETMAIN. (When a job issues a variable length GETMAIN, the amount of storage obtained is the difference between REGION ASSIGNED and the top of the allocated area, assuming that the largest free block is there.)

The information reported under PRIVATE STORAGE MAP defines significant boundaries within the private area and shows the space between them. There is a separate map for storage below the 16-megabyte line and for extended storage (above the 16-megabyte line). Each map identifies:

Figure 3-43. PRIVATE STORAGE MAP Heading in Private Area Summary Report

Field Heading	Definition
BOTTOM OF ALLOCATED AREA	The lowest address of allocated storage for LSQA, SWA, and subpools 229 and 230, all of which are allocated down from the top of the private area. RMF reports the lowest value it found during the RMF interval. The time stamp reflects the time when RMF first observed the value reported.
GETMAIN LIMIT	For virtual storage below and above the 16-megabyte line, shows the installation limit on the total amount of storage a job can obtain via GETMAIN macro instructions. The values reported are set by your installation (using the IEFUSI installation exit or IEALIMIT) or the system defaults. If a job requests storage that the system would have to obtain from storage above either of these limits, the job terminates abnormally.
TOP OF ALLOCATED AREA	The highest address of user region storage allocated up from the bottom of the private area, including subpools 251 and 252 as well as user subpools 1 through 127. RMF reports the highest value it found during the RMF interval. The time stamp reflects the time when RMF first observed the value reported.

It is possible for the top and bottom of the allocated area to cross. When RMF detects this situation, it reports a negative value for the area between the marks; it does not reverse the labels.

The bottom half of the report contains information about free and allocated storage within the private area, both below and above the 16-megabyte line. For LSQA/SWA/229/230 and for USER REGION, RMF reports the following information:

Figure 3-44 (Page 1 of 2). Private Area Summary Report (Bottom Half)

Field Heading	Definition
FREE PAGES (BYTES)	Presents the MIN, MAX, and AVG values, in bytes, for the amount of free storage.
LARGEST FREE BLOCK	Presents MIN, MAX, and AVG values, in bytes, for the size of the largest free block in the available amount of free storage.

Figure 3-44 (Page 2 of 2). Private Area Summary Report (Bottom Half)	
Field Heading	**Definition**
PAGES ALLOCATED (IN BYTES)	Presents MIN, MAX, and AVG values for the amount of allocated storage. **Note:** The MAX value for PAGES ALLOCATED has the same meaning as fields in SMF record type 30. For LSQA, these fields are SMF30ARB and SMF30EAR. For the user region, these fields are SMF30URB and SMF30EUR. RMF, however, might report a smaller number than SMF does because RMF reports the highest value that it sampled while SMF reports the highest value that occurred. If the highest value occurred when RMF was not taking a sample, RMF misses the actual peak value.

For USER REGION and LSQA, RMF determines FREE PAGES and LARGEST FREE BLOCK in relation to the GETMAIN limit.

It is possible that LSQA/SWA can become so large that it extends below the GETMAIN limit. Thus, a GETMAIN macro instruction for user region storage fails even though the storage requested does not exceed the GETMAIN limit. In this case, RMF makes the appropriate adjustments to the values it reports for FREE PAGES and LARGEST FREE BLOCK.

Partial Private Area Data: Private area reporting works best for long-running non-swappable jobs, such as JES2, JES3, VTAM, or IMS. You can, of course, monitor other jobs, but there are some conditions that mean RMF can collect little or no data. These conditions, which are related to the way virtual storage reporting works, are:

1. RMF searches for any requested jobs at the beginning of each interval. If it does not find a job, it does not monitor the job during the interval. In this case, RMF issues a message to the operator and produces a report. The report, however, contains no data; instead, the following message appears:

   ```
   JOB WAS NOT ACTIVE AT THE BEGINNING OF THIS INTERVAL
   ```

   RMF continues to search for the job at the beginning of each interval. When it finds the job, it deletes the message, monitors the job, and produces a report.

   If a job begins and ends within a single RMF interval, RMF cannot monitor its use of virtual storage.

2. If a job that RMF is monitoring terminates and is then restarted, the report for the interval in which it terminated includes data only up to the point when the job terminated. RMF resumes its monitoring of the restarted job at the beginning of the interval following the interval during which the job was restarted.

3. If a job RMF is monitoring is swapped out at the time RMF takes a sample of virtual storage data, RMF does not cause a swap-in; it skips the sample for that job. Thus, the number of samples for a swappable job may be less than expected. If a job is swapped out every time RMF tries to take a sample during an interval, RMF reports no data for that interval.

Private Area Detail Report

The optional private area detail report contains additional information about the use of allocated storage in the private area below the 16-megabyte line. It includes MIN, MAX, and AVG values for the number of bytes of allocated blocks, broken down by subpool and by area (LSQA, SWA).

Figure 3-45 shows a sample printed report for the private area (detail).

```
                                 V I R T U A L   S T O R A G E   A C T I V I T Y
                                                                                          PAGE 1

       OS/VS2                  SYSTEM ID AQFT              DATE 03/04/89            INTERVAL 29.59.997
       SP2.2.0                 RPT VERSION 3.5.0           TIME 09.00.00            CYCLE 1.000 SECONDS

                                         PRIVATE AREA DETAIL
  JOB NAME -  VTAM
  NUMBER OF BYTES OF ALLOCATED BLOCKS BY AREA (BELOW 16 MEG)
  SUBPOOL (AREA)       MIN            MAX            AVG
    229                32K 09.00.01   32K 09.00.01   32K
    230                80K 09.00.01   80K 09.00.01   80K
    236 (SWA)          64K 09.00.01   64K 09.00.01   64K
    237 (SWA)          12K 09.00.01   12K 09.00.01   12K
    255 (LSQA)         36K 09.00.01   36K 09.00.01   36K
  USER REGION
      0                 4K 09.00.01    4K 09.00.01    4K
    251 (MODULES)      36K 09.00.01   36K 09.00.01   36K
    252 (REENTRANT)     4K 09.00.01    4K 09.00.01    4K
```

Figure 3-45. Private Area Detail Report

Changing the Sample to Cycles Ratio

To minimize its overhead, RMF collects one sample of virtual storage data for every ten RMF cycles. Expressed as a ratio, the RMF default is 1:10. Your installation might want to change this default ratio.

Decreasing the 1:10 ratio -- for example, to 1:5 -- increases the accuracy of the virtual storage data RMF collects, in the sense that RMF is more likely to capture such data as a peak value when it samples more frequently. Decreasing the ratio, however, does increase RMF overhead.

Increasing the 1:10 ratio -- for example, to 1:15 -- decreases the accuracy of the virtual storage data RMF collects, in the sense that RMF is less likely to capture such data as a peak value when it samples less frequently. Increasing the ratio, however, does decrease RMF overhead.

To change the ratio, use the AMASPZAP program. For example, to change the ratio from 1:10 to 1:15 -- so that RMF takes one virtual storage sample for every fifteen cycles -- use the following JCL and control statements:

```
//ZAP JOB
//STEP EXEC PGM=IMASPZAP
//SYSPRINT DD SYSOUT=A
//SYSLIB DD DSN=SYS1.LPALIB,DISP=SHR
//SYSIN DD *
  NAME ERBMFEVS NUMCYCLE
  VERIFY 00 0000000A
  REPL   00 0000000F
/*
```

Note that the change does not take effect until after the operator performs the next cold start (IPL with CLPA) of the system.

For more information on the use of AMASPZAP, see *MVS/XA Service Aids*, GC28-1159.

Workload Activity Report

The workload activity report provides information about the utilization of each performance group period (PGP) within a performance group. Each performance group can be defined in the installation control specification as well as in the IPS.

To enable you to obtain the information in the format best suited to your needs, nine types of sub-reports are available. The nine ways in which the data can be formatted are:

1. **By performance group period within performance group** - the most detailed report. Gives information on each performance group period. A summary for each performance group immediately follows the information for all performance group periods in the performance group; however, this summary is not produced if all performance group periods have zero interval service.
2. **By performance group** - summarizes the data for all performance group periods within the performance group.
3. **By performance objective in a range of performance groups** - summarizes the data by performance objective within a range of ten performance groups.
4. **By domain in a range of performance groups** - summarizes the data by domain number within a range of ten performance group numbers.
5. **By performance group range** - summarizes the data within a range of ten performance group numbers.
6. **By performance objective for entire system** - summarizes the data by performance objective for the entire system.
7. **By domain for entire system** - summarizes the data by domain for the entire system.
8. **For entire system** - summarizes the data for the entire system.
9. **By time slice group** - summarizes the data for all performance group periods within the time slice group. This report can be produced only when time slice groups have been defined in the IPS.

Figure 3-46 shows the options that can be specified when you request workload activity and the sub-report types that are produced for each option.

Option	1	2	3	4	5	6	7	8	9
PERIOD	X								
GROUP		X							
RANGE			X	X	X				
OBJ						X			
DOMAIN							X		
SYSTEM								X	
TIME									X

Figure 3-46. Sub-Report Types Included by Option

You can assign unique performance group numbers and domain numbers to a particular user or set of users. This enables you to use the performance group period report to obtain individual workload reporting for users such as subsystems, monitor programs, groups of users, or batch initiator classes.

The information given in the service definition coefficients and the IOC, CPU, MSO, and SRB interval service fields can be used to determine the actual number of countable EXCP instructions, the task processor (CPU) time, the approximate storage use for each performance group, and the SRB processor time. The basic formula for determining the processor, I/O, SRB, or storage service is:

$$\left\{\begin{matrix}\text{CPU}\\\text{IO}\\\text{MSO}\\\text{SRB}\end{matrix}\right\} \text{ Service Units } = \frac{\left\{\begin{matrix}\text{CPU}\\\text{IOC}\\\text{MSO}\\\text{SRB}\end{matrix}\right\} \text{ service}}{\left\{\begin{matrix}\text{CPU}\\\text{IOC}\\\text{MSO}\\\text{SRB}\end{matrix}\right\} \text{ service definition coefficient}}$$

See *MVS/XA Initialization and Tuning Guide,* GC28-1149, for more details.

The number of swaps in the workload activity report can be cross-checked with the paging report swap-out counts to point out the group period where the swap activity is occurring. This information is useful when you are trying to eliminate unnecessary swapping.

Note, however, that the swap count in the workload activity report does not include any swaps that occur prior to the INITATT sysevent being issued or any swaps that occur after the INITDET sysevent is issued.

The workload activity report provides information about the utilization of each PGP defined in the IPS (IEAIPSxx parmlib member) and the installation control specification (IEAICSxx parmlib member). The definition in the installation control specification for a performance group can include a subsystem (SUBSYS), transaction class (TRXCLASS), account information (ACCTINFO), user identifier (USERID), and transaction name (TRXNAME). To enable you to associate data for a performance group with the definition for that performance group in the installation control specification, RMF includes this information in the type 1 and type 2 sub-reports (the reports by performance group period and by performance groups).

The system resources manager (SRM) uses the values in the IPS and the installation control specification to perform its workload balancing functions. When the SET operator command is used to change the SRM values contained in the IPS or the installation control specification when RMF is measuring workload activity, the SRM notifies RMF. RMF then performs normal end-of-interval processing for all active measurements and reinitializes all active measurements. That is, RMF uses the old SRM values to generate SMF records and, if REPORT is in effect, reports for the part of the interval that had elapsed before the SRM values changed. RMF then begins a new interval using the new SRM values.

The installation control specification also enables you to define performance groups designed solely to report on subsystem transaction activity. These report performance groups are included only in the type 1 and type 2 sub-reports (the reports by performance group period and by performance group). Report performance groups have no performance objective, domain, or time slice group. Therefore, the OBJ NUMBER, DOMAIN NUMBER, and TIME SLICE GROUP fields for a report performance group period in the type 1 subreport contain two asterisks(**). Also, the value reported for WORKLOAD LEVEL for a report performance group is always zero.

The other fields in the report can contain, for report performance groups, only the data that a subsystem provides to the SRM. The system-defined subsystems (TSO and STC) and the job entry subsystems (such as JES2 or JES3) provide meaningful data for all fields in the report, except those noted earlier as not applying to report performance groups. For other subsystems (such as a user-written subsystem), only ENDED TRANSACTIONS and AVG TRANS TIME contain meaningful data, and INTERVAL SERVICE might contain meaningful data. (See the current documentation for the subsystem to determine if it provides service information to the SRM.) When data for a field is not meaningful, the field contains zeroes.

Note: A TSO user can press the TEST REQ key on a 3277 terminal to suspend the TSO session and then resume the session by a LOGON reconnect command. For the time that elapses between these two actions, the system places the user in a wait state and the SRM continues to accumulate CPU and transaction time for the user. If RMF is producing workload activity reports, the time is included in the AVG TIME OF ENDED TRANS (TSO response time) thus, the time reported might be greater than the actual average time for normal transactions.

Data Fields

In addition to the report headers that are common to all reports, the workload activity report header contains:

- the IPS identification number
- the sub-report title (one of the nine types listed above)
- the OPT parmlib member names and the installation control specification
- the values for the IOC, CPU, MSO, and SRB service definition coefficients as defined in the IPS specification
- for a 3090 processor, the number of SRM processor service units per CPU second (SU/SEC=xxx). This heading does not appear in reports for 308X / 908X or 4381 processors.

Duration Report: When you request a workload activity duration report, the message IPS CHANGED DURING DURATION INTERVAL will appear at the end of every sub-report if the IPS changed during the duration of time covered in the report. The message indicates a change to the SRM values in the IPS, installation control specification, or OPT. Because the SRM informs RMF when such a change occurs, the values reported are meaningful.

Performance

The type 1 and type 2 sub-reports (the reports by performance group period and by performance group) can include two lines of information that precede the data for each performance group. These lines identify the parameters in the installation control specification entry for the performance group.

These lines appear only when the installation control specification (ICS) associates the performance group with a single subsystem name.

They do not appear in the report for any one of the following reasons:

- If more than one subsystem is associated with the performance group.
- If no subsystem is associated with the performance group.
- If there is no active installation control specification.

When the lines appear, the subsystem name is always reported, and the lines have the format:

```
SUBSYS=subsystem name TRXCLASS=transaction class ACCTINFO=account information
USERID=user identifier TRXNAME=transaction name
```

The report lists user identifier, transaction class, transaction name or accounting information only when (1) the corresponding parameter appears in the installation control specification entry for the performance group and (2) the parameter is the only user identifier, transaction class, or transaction name associated with the performance group.

Figure 3-47 indicates how the various data fields are presented for the different sub-report types.

Figure 3-48 and Figure 3-49 illustrate the type 1 and type 8 workload activity reports.

	Sub-Report Type							
Data Field	1	2	3	4	5	6	7	8
PERFORMANCE GROUP	E	E	B	B	B	C	C	C
GROUP PERIOD	E	C	C	C	C	C	C	C
OBJECT NUMBER	D	C	E	C	C	E	C	C
DOMAIN NUMBER	D	C	C	E	C	C	E	C
TIME SLICE GROUP	D	C	C	C	C	C	C	C
INTERVAL SERVICE	A	A	A	A	A	A	A	A
AVERAGE ABSORPTION RATE	A	A	A	A	A	A	A	A
AVERAGE TRANS SERV RATE	A	A	A	A	A	A	A	A
WORKLOAD LEVEL	A	A	A	A	A	A	A	A
PAGE IN RATE	A	A	A	A	A	A	A	A
STORAGE AVERAGE	A	A	A	A	A	A	A	A
TOTAL	A	A	A	A	A	A	A	A
AVERAGE TRANSACTIONS	A	A	A	A	A	A	A	A
MPL	A	A	A	A	A	A	A	A
ENDED TRANSACTIONS	A	A	A	A	A	A	A	A
NUMBER OF SWAPS	A	A	A	A	A	A	A	A
AVG TIME OF ENDED TRANS	A	A	A	A	A	A	A	A

A The data is provided unless all data for the entry in zero; in this case 'ZEROS' appears in the service in interval area and all other entries are blank.
B A range of ten performance groups is given.
C 'ALL' is printed rather than a specific number.
D A value is printed when it is unique; when the value is not unique, 'ALL' is printed.
E A range of values can be produced if all performance group periods that make up the range have interval service.

Figure 3-47. Data Fields Per Sub-Report Type

```
                                          W O R K L O A D   A C T I V I T Y
                                                                                                            PAGE 1

          OS/VS2                      SYSTEM ID AQFT          DATE 03/04/89            INTERVAL 29.59.997
          SP2.2.3                     RPT VERSION 3.5.1       TIME 09.00.00            IPS = IEAIPSFT

        OPT = IEAOPTFT                         REPORT BY PERFORMANCE GROUP    SERVICE DEFINITION COEFFICIENTS     SU/SEC= 625.9
        ICS = IEAICSFT                                    PERIOD              IOC =  5.0  CPU = 10.0  MSO =  3.0  SRB = 10.0

 **** PERFORMANCE ***** DOMAIN  TIME   INTERVAL SERVICE  AVERAGE ABSORPTION,   PAGE     STORAGE    AVERAGE    ENDED      AVG TRANS
 GROUP    GROUP  OBJTVE NUMBER  SLICE  (TOTAL, BY TYPE   AVG TRX SERV RATE,    IN       AVERAGE,   TRANS,     TRANS,     TIME, STD
 NUMBER   PERIOD NUMBER         GROUP  AND PER SECOND)   WORKLOAD LEVEL        RATE     TOTAL      MPL        #SWAPS
HHH.MM.SS.TTT
 0000       1     00     000     **    IOC=    18,961    ABSRPTN =    1,980     .00      603.96      4.34          1
006.27.04.922
                                       CPU= 1,384,730    TRX SERV=    1,979
                                       MSO= 13.730448M   WKLD LEV=   148.66              2624.9      4.34        124
000.00.00.000
                                       SRB=   355,823
                                       TOT= 15.489962M
                                       PER SEC=   8,605

 SUBSYS = JES2        TRXCLASS =              ACCTINFO = NO
 USERID =              TRXNAME =
00001       1     06     001     **    IOC=   732,711    ABSRPTN =    1,812     .00      240.44      1.46         57
000.00.13.470
                                       CPU= 1,453,698    TRX SERV=    1,784
                                       MSO= 2,395,910    WKLD LEV=    44.56              345.78      1.43        116
000.00.23.750
                                       SRB=   110,311
                                       TOT= 4,692,630
                                       PER SEC=   2,607
00001       2     01     002     **    IOC= 1,778,011    ABSRPTN =    4,445     .00      410.00      3.51         45
000.04.29.495
                                       CPU= 7,221,286    TRX SERV=    4,432
                                       MSO= 18.881666M   WKLD LEV=   136.95              1437.8      3.50         50
000.13.21.600
                                       SRB=   179,794
                                       TOT= 28.060757M
                                       PER SEC=  15,589
00001      ALL   ALL     ALL     ALL   IOC= 2,510,722    ABSRPTN =    3,679     .00      360.68      4.97        102
003.02.06.422
                                       CPU= 8,674,984    TRX SERV=    3,655
                                       MSO= 21.277576M   WKLD LEV=   123.71              1783.6      4.94        166
000.09.04.600
                                       SRB=   290,105
                                       TOT= 32.753387M
                                       PER SEC=  18,196

 SUBSYS = TSO         TRXCLASS =              ACCTINFO = NO
 USERID =              TRXNAME =
00002       1     01     003     **    IOC=   527,605    ABSRPTN =    2,312     .00      260.18      1.34      14065
000.00.00.079
                                       CPU= 2,225,935    TRX SERV=    2,287
                                       MSO= 2,499,198    WKLD LEV=   126.90              344.89      1.32      13929
000.00.00.479
                                       SRB=   265,046
                                       TOT= 5,517,784
                                       PER SEC=   3,065
```

Figure 3-48. Workload Activity Report by Performance Group Period


```
                                        W O R K L O A D   A C T I V I T Y
                                                                                                            PAGE 1

             OS/VS2                    SYSTEM ID AQFT           DATE 03/04/89            INTERVAL 29.59.997
             SP2.2.3                   RPT VERSION 3.5.1        TIME 09.00.00            IPS = IEAIPSFT

          OPT = IEAOPTFT                              SYSTEM SUMMARY          SERVICE DEFINITION COEFFICIENTS      SU/SEC= 625.9
          ICS = IEAICSFT                                                      IOC =  5.0  CPU = 10.0  MSO =  3.0  SRB = 10.0

 **** PERFORMANCE ***** DOMAIN  TIME   INTERVAL SERVICE  AVERAGE ABSORPTION,   PAGE     STORAGE    AVERAGE    ENDED     AVG TRANS
 GROUP    GROUP  OBJTVE NUMBER  SLICE  (TOTAL, BY TYPE   AVG TRX SERV RATE,    IN       AVERAGE,   TRANS,     TRANS,    TIME, STD
DEV
 NUMBER   PERIOD NUMBER         GROUP  AND PER SECOND)   WORKLOAD LEVEL        RATE     TOTAL      MPL        #SWAPS
HHH.MM.SS.TTT
 ALL         ALL  ALL    ALL    ALL    IOC= 6,710,407    ABSRPTN =    1,526     .00      442.14       33.99    18900
000.00.02.538
                                       CPU= 25.504230M   TRX SERV=    1,523
                                       MSO= 59.221280M   WKLD LEV=  117.12               15001.       33.92    19385
000.02.53.800
                                       SRB= 1,795,025
                                       TOT= 93.230942M
                                       PER SEC=  51,795
```

Figure 3-49. Workload Activity System Summary

Figure 3-50 (Page 1 of 3). Fields in the Workload Activity Report	
Field Heading	**Definition**
GROUP NUMBER	The identification of the performance group that describes the rate at which service is to be allocated to the user. The report can show a range of PG numbers to eliminate printing of consecutive lines of zeros. The range of values is 0 to 9999.
GROUP PERIOD	The identification of the performance group period that is a portion of the elapsed transaction time, such as the average service used by a transaction during the first 10 seconds of a total transaction time of 30 seconds. The report can show a range of PGP numbers to eliminate printing of consecutive lines of zeros. The range of values is 1 to 8.
OBJTVE NUMBER	The identification of the performance objective associated with a performance group. The range of values is 0 to 64. If no performance objective was defined for the performance group period, two asterisks (**) appear in this field.
DOMAIN NUMBER	The identification of the domain associated with the performance group during a performance group period. The range of values is 0 to 128. If no domain number was defined for the performance group period, two asterisks (**) appear in this field.
TIME SLICE GROUP	The identification of the time slice group associated with the performance group period. The range of values is from 1 to 16. Two asterisks (**) in this field indicate that no time slice group has been defined for the performance group period.

Figure 3-50 (Page 2 of 3). Fields in the Workload Activity Report

Field Heading	Definition
INTERVAL SERVICE	This category is made up of five distinct data values: **IOC** = The total amount of input/output service received. **For type 1 sub-reports**, this is the I/O service received in one period. **For all other types**, it is the sum of the I/O service received by all periods being summarized. **CPU** = The total amount of TCB processor service received. **For type 1 sub-reports**, this is the service received in one period. **For all other types**, it is the sum of the service received by all periods being summarized. **MSO** = The total amount of main storage service received. **For type 1 sub-reports**, this is the main storage service received in one period. **For all other types**, it is the main storage service received by all periods being summarized. **SRB** = The total amount of SRB service received. **For type 1 sub-reports**, this is the SRB service received in one period. **For all other types**, it is the SRB service received by all periods being summarized. **TOT** = The total amount of service received. **For type 1 sub-reports**, this is the total service received in one period. **For all other types**, it is the sum of the total service received by all periods being summarized. **PER SEC** = The rate at which service is provided in units of service per second. **For type 1 sub-reports**, this is the total period service divided by the RMF interval time. **For all other types**, it is the sum of the total service for all periods being summarized, divided by the RMF interval time.
AVERAGE ABSORPTION, AVG TRX SERV RATE, WORKLOAD LEVEL	This category is made up of three distinct data values: **ABSRPTN** = The rate at which service is used while transactions are resident in real storage. **For type 1 sub-reports**, this is the total period divided by the transaction active time in storage in the period. **For all other types**, it is the sum of the total service in the periods being summarized, divided by the sum of the transaction active time in storage for all the periods being summarized. **TRX SERV** = The rate at which service is used by transactions that are active but not necessarily in storage. **For type 1 sub-reports**, this is the total period service, divided by the transaction active time in the period. **For all other types**, it is the sum of the total service in the periods being summarized, divided by the sum of the transaction active time for all the periods being summarized. **WKLD LEV** = The average SRM recommendation value given to users based on service usage and IPS objective specifications. **For type 1 sub-reports**, it is the value of the workload level for the period. **For all other types**, it is the sum of the products of the workload level and the total service for each period being summarized, divided by the sum of the total service for each period being summarized.
PAGE IN RATE	The average rate at which pages are read into real storage while transactions are resident in real storage. **For type 1 sub-reports**, this is the total number of page-ins during the period, divided by transaction residency time. **For all other types**, it is the sum of the total number of page-ins for all periods that are summarized, divided by the sum of the transaction residency time for all the periods being summarized.
STORAGE AVERAGE, TOTAL	This heading defines two distinct data fields in the same column. The first row contains the average storage, followed by a skipped line. The following row contains the total storage value. **Average Storage:** The weighted average number of frames, real and extended, allocated to active ASIDs. **For type 1 sub-reports**, this value is the sum of the number of real and extended frames weighted by the transaction residency time for each active ASID, divided by the total transaction residency time for the period. **For all other types**, this value is the sum of the number of real and extended frames weighted by the residency time, divided by the total of all the transaction residency times for all the periods being summarized. **Total Average:** The total number of real and extended storage frames allocated to resident ASIDs. **For type 1 sub-reports**, this value is the sum of the total number of real and extended storage frames weighted by the transaction residency time for each active ASID, divided by the RMF interval time. **For all other types**, it is the sum of the total number of real and extended storage frames weighted by the transaction residency time of each active ASID, divided by the RMF interval time, for all the periods being summarized.

Figure 3-50 (Page 3 of 3). Fields in the Workload Activity Report

Field Heading	Definition
AVERAGE TRANSACTIONS, MPL	This heading defines two distinct data fields in the same column. The first row contains the average transaction field followed by a skipped line. The following row contains the multi-programming level (MPL) value. **Average Transactions:** The average number of active transactions during the interval. **For type 1 sub-reports,** this is the transaction active time for the period, divided by the RMF interval time. **For all other types**, it is the sum of the transaction active time for all summarized periods, divided by the RMF interval time. **MPL:** The average number of transactions resident in real storage during the interval. **For type 1 sub-reports,** the multi-programming level is the transaction residency time for the period, divided by the RMF interval time. **For all other types**, it is the sum of the transaction residency time for all summarized periods, divided by the RMF interval time.
ENDED TRANS, #SWAPS	This heading defines two distinct data fields in the same column. The first row contains the ended transactions field followed by a skipped line. The following row contains the number of swaps. **Ended Transactions:** The number of transactions that ended during the interval. **For type 1 sub-reports**, this is the number of transactions that ended during the period. **For all other types**, it is the total number of transactions that ended for all the summarized periods. **# Swaps:** The total number of swaps. **For type 1 sub-reports**, it is the number of swaps that occurred during the period. **For all other types**, it is the sum of the number of swaps that occurred during all summarized periods.
AVG TRANS TIME / STD DEV (HH.MM.SS.TTT)	The average elapsed time of ended transactions. **For type 1 sub-reports**, this is the elapsed time of ended transactions for the period, divided by the number of ended transactions for the period. **For all other types**, it is the sum of the elapsed time of ended transactions for all periods being summarized, divided by the sum of the number of ended transactions for all periods being summarized. For TSO and other subsystems, this field represents the response time. The average elapsed time is followed by the standard deviation. Standard deviation is expressed in the same format as the average elapsed time.

Chapter 4. Monitor II Session

There are two types of Monitor II sessions: the Monitor II display session and the Monitor II background session. These sessions vary in the form of output produced. The two sessions are described under "Monitor II Display Session Control" and "Monitor II Background Session Control." "Monitor II Session Reports" describes each report and gives report examples.

Monitor II Display Session Control

There are two types of Monitor II display sessions: the local 3270 display session and the TSO display session. The major difference between them occurs when they are started and ended.

Starting and ending each type of session is described under "Starting and Ending a TSO Display Session" and "Starting and Ending a Local 3270 Display Session."

Session processing for both types of display sessions is described under "Controlling Display Session Processing," and a list of the valid menu items appears under "Menu Items."

Starting and Ending a TSO Display Session

A Monitor II TSO display session differs from other RMF sessions in that, unless you need reports based on data the Monitor I session collects, it is independent of RMF initialization.

To start a TSO session, logon to TSO. When the logon procedure is complete, and TSO is ready to receive your commands, enter the command

```
RMFMON
```

to invoke RMF. RMFMON can also be entered at any point during a TSO session. The response to RMFMON builds the control blocks your session requires, creates a display session task, and readies the Monitor II display functions of RMF. The response might also include calling a routine supplied by your installation to verify that you are authorized to use RMF. RMF displays the menu when it is ready to receive your display session commands. TSO appears in the session-id field of all RMF messages issued to you during the RMF session.

Session processing is described under "Controlling Display Session Processing."

To end the RMF session, use the stop display command, which is Z. You can then either logoff TSO or continue with other TSO work.

Starting and Ending a Local 3270 Display Session

The system operator must start a Monitor II local 3270 display session. The operator or the terminal user can end it.

Starting the Session

An RMF START session command starts the session; the command is passed to RMF through the input field of the system MODIFY command. The command is issued from the operator console. The required syntax is:

```
{MODIFY} {[RMF.]id,} {START} nnn
{  F   } {jobname, } {  S  }
```

id
The identifier assigned to RMF when it was initialized.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

nnn
The three digit hexadecimal device number of the display station to be used for the session. This device number is the identifier for your session; it appears in all RMF messages issued to you. It must be unique; it cannot be the same as an identifier assigned to another currently active session. Specifying a session identifier that is not unique or is not the valid device number of a display station will cause an error message.

When a valid START command is issued, RMF responds by building the control blocks your session requires, creating a display session task, and readying the Monitor II display functions of RMF. A message is issued to the operator console when RMF is ready to receive your display commands. The menu is displayed on the screen.

Example of Starting the Session

To start a Monitor II local 3270 display session at a display station with a device number of 642, issue the command:

```
MODIFY RMF.A,START 642
```

Ending the Session

The terminal user or the system operator can end the session. When a session is ended, RMF issues a message to the operator console and terminates the session.

The terminal user can end the session by entering the stop display command, which is Z.

The system operator can end the session by issuing either a STOP system command, which ends all active non-TSO RMF sessions, or a STOP session command for your particular session. The STOP session command is passed to RMF through the input field of the system MODIFY command. The syntax required is:

```
{MODIFY} {[RMF.]id,} {STOP} nnn
{  F   } { jobname,} { P  }
```

id
The identifier assigned to RMF when it was initialized.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

nnn
The three-digit device number session identifier assigned when your session was started.

Example of Ending the Session

To end a session started with a session identifier of 642, issue the command:

```
MODIFY RMF.A,STOP 642
```

Controlling Display Session Processing

When RMF is ready to receive your commands, the menu frame appears on the display screen. The format of the display screen for an RMF display session, the use of the display station keys, and the display commands you can issue to RMF are described in the following sections.

Display Screen Format

Figure 4-1 shows the format of the screen for an RMF display session. The use of each field is described in the following text. The letters in the following figure refer to the letters in the text.

(A) The input area. Use this field to enter display commands, such as a request for a report. At the beginning of the session or after RMF has responded to a display command, the cursor is positioned at the beginning of the field. The input area is the only unprotected area of the screen; that is, it is the only area whose contents you can modify. For those systems running on a 3090 processor, migration age (the amount of time a page resides on extended storage before it migrates to auxiliary storage) appears in the input area of the display screen.

The input area is blank or, for systems running on a 3090 processor, contains only MIG=xxx at the beginning of the session and after RMF has responded to a command, unless one of the following situations has occurred:

1. You have entered a display command that contains a syntax error. In this case, the command is displayed in the input area to enable you to correct and re-enter the command. A message describing the error appears in the status/message area (area **(B)** in the figure).

   When you are working on a 3090 processor and have entered a command in the input area, at least one blank must separate an input command or operand from the displayed MIG=xxx, or the input must use all 32 positions. Otherwise, RMF treats MIG=xxx as part of the input and thus detects a syntax error; it redisplays the invalid input and does not update the migration age. To ensure that this does not happen, leave one blank in the command line after you enter a command or operand that does not use all 32 positions.

2. You have entered a timed update command (T). RMF uses the input area to display the length of the time interval currently being measured and the number of intervals remaining before the timed update ends.

   If your system runs on a 3090 processor, the message TIMED SCREEN UPDATE that appears on the screen is shortened to TIME UPDATE to avoid overlaying the migration age (MIG=xxx).

3. You have requested a table report that is longer than the maximum number of lines for the display screen. The maximum number of lines for each type of display screen is:

   - 21 lines for a 3277 Model 2, 3278 Model 2, 3276 Model 2, or a 3279 Model 2A or 2B
   - 29 lines for a 3276 Model 3, 3278 Model 3, or a 3279 Model 3A or 3B
   - 40 lines for both a 3276 Model 4 and a 3278 Model 4

   For example, if you are using a 3277 Model 2, an F appears in the input area (area **(A)** in the figure) and the first 21 lines of data appear in the output area (area **(F)** in the figure). The F indicates that you should enter the frame command to view the next frame of output data. Because the area contains an F, you can view the next frame by pressing the enter key. The F continues to appear in the input area until all frames of data have been viewed. When the last frame is displayed, the end of the report is indicated by a blank input area. Should you decide at any point that you do not need to see the remaining frames, blank out the F in the input area and press the enter key, or enter your next command.

(B) The status/message area.

When you have entered an invalid command or a command with a syntax error, a message describing the error appears in the status/message area. When you have entered a timed update (T) command, a message requesting you to enter the report options appears in the status/message area. It normally contains the following general system load indicators:

CPU - the current processor utilization This field is set to 101% when the actual utilization is 100% and one or more users are waiting to be dispatched.

UIC - the highest unreferenced interval count

PDT - the current SRM view of the page delay time (that is, the number of milliseconds required to process a paging request)

DPR - the current SRM view of the demand paging rate (that is, the sum of the number of non-swap, non-VIO page-ins per second and the number of non-swap, non-VIO page outs per second.)

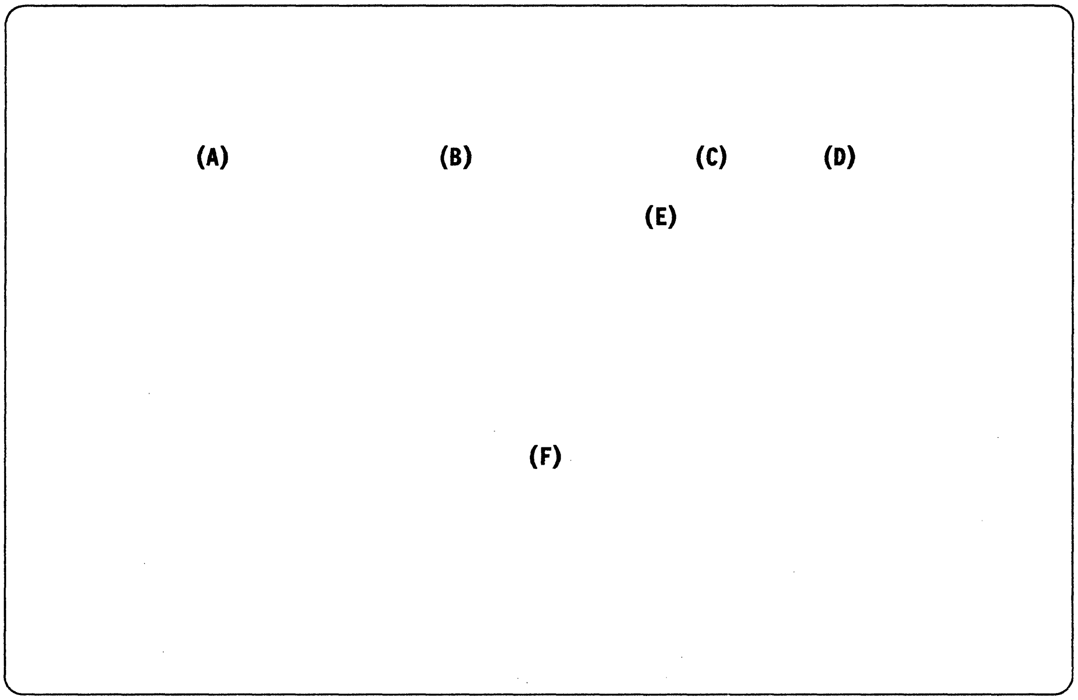

Figure 4-1. Display Station Screen Format

(C) The report name field. It contains the name of the report whose output is currently shown in the output area (area **(F)** in the figure). If there is no report output shown, this field contains the name of the report that will be generated if you press the enter key without entering a display command. The report name field is blank when the session begins. If you press the enter key when the report name field is blank at the beginning of the session, the first report in the default menu will be produced, using default operands.

Note: If you change the first entry in the menu by means of a set PF key operation while the report field is blank, the PF key setting is changed. If, however, you press the enter key while the report name field is blank, the first report listed in the original default menu is produced.

(D) The mode field. It indicates the current setting of the delta and hardcopy modes for the session. The field contains two characters.

The first character can be either D, indicating delta mode for the session, or T, indicating total mode for the session. See "Set Delta Mode" later in this section for more information. D appears after you have entered a D ON or D command; T appears at the beginning of the session or after you have entered a D OFF command.

The second character can be either an H, indicating that hardcopy mode is in effect, or a blank, indicating that hardcopy mode has not been specified. See "Set Hardcopy Mode" later in this section for more information. H appears after you have entered an H ON or H command. A blank appears at the beginning of the session or after you have entered an H OFF command.

(E) The header area. When the output area (area **(F)** in Figure 4-1) contains report data, the two lines of the header area contain the column headings for the displayed report. When a processing error occurs, descriptive messages appear in this area, and no data is displayed in the output area.

The output area. It contains report output displayed in response to your report request. At the beginning of the session or in response to a menu (M) command, the menu frame appears in the output area. Occasionally, the output area can contain an informational message, such as a message which is issued when you have requested a report for which no data is currently available.

Display Station Keys

The only keys that can be used to transmit data from the display devices to RMF are the enter key and selected program function keyboard (PF keys). (For a list of the devices RMF supports, see "System Requirements" in Chapter 1.) Normally, you would enter your commands in the input area and press the enter key to transmit the command to RMF. However, there are three situations when you can press the enter key without previously entering a command. These situations are:

1. At the beginning of a session when the menu frame is displayed in the output area, you can press the enter key without entering a command. This will cause the first report in the menu to be generated, with all operands defaulted.

   Note: If you change the first entry in the menu by means of a set PF key operation while the report field is blank, the PF key setting is changed. If, however, you press the enter key while the report name field is blank, the first report listed in the original default menu is produced.

2. If, after entering a report request and inspecting the output, you want to make a repetitive request for the same report, press the enter key without respecifying the report request. Pressing the enter key in this situation will cause the previously requested report to be generated again. The operands RMF uses are the same as those specified on the original request. This use of the enter key enables you to build a table of repetitive requests with a minimum of input and is especially useful when you are requesting a row report.

3. If an F appears in the input area, indicating that there are subsequent frames in the table report you have requested, you can view the next frame by pressing the enter key.

For a TSO session, the attention key (PA1) can be used in the same way it can be used for any TSO command as long as the timed update command is not active. The timed update command is described later in this section under "Repeating Reports." For a local 3270 session, pressing the attention key causes RMF to display the menu frame. Pressing PA2 causes RMF to re-display the the previous screen and wait for input. The program function keys (PF keys) can be used to request RMF reports as described later in this section under "Set PF Keys."

Display Commands

Display commands enable you to perform the following functions:

- display the menu frame at any time
- display the default operands associated with each menu item
- override the PF keys
- set the delta mode or total mode for the session
- set the hardcopy mode for the session
- request a report
- select an item from the menu and cause the display of the selected report to be updated automatically at a timed interval
- frame through a table report
- recall the previous data
- print a report
- stop the session

The functions and the commands that they require are described in the following sections.

Display the Menu

To display the menu, issue the menu command. The syntax is:

```
M
```

When you enter M, RMF responds by displaying the menu of reports available. Figure 4-2 shows the menu frame. The menu frame lists each menu item, the PF key associated with the menu item, a description of the corresponding report, and the IPS parmlib member name. If your installation has added additional reports to those provided by RMF, the menu item and, possibly, a description of the report also appear in the menu frame.

Display Default Operands

To display the default operands associated with each menu item, issue the MM menu command. The syntax is:

```
MM
```

When you enter MM, RMF responds by displaying a menu frame that contains the report name, the current PF key assignments, and the default operands for both the data gatherer and data reporter.

The data gatherer and reporter options are separated by three dashes (---). If both of the default options are the same, only one set of options is displayed. If there are no default options for the data gatherer, the reporter options are displayed preceded by three dashes (---).

You can change the PF key assignments by means of the PF key keyword of the PICTURE macro or by means of the # command described next under "Set PF Keys." The PICTURE macro is described in Chapter 7 under "Using the PICTURE Macro."

Set PF Keys

RMF assigns only the first 12 program function keyboard (PF keys) as shown in Figure 4-2. You can override any RMF-assigned key for the duration of your Monitor II session by entering a pound sign (#) in the data input area, followed immediately by a report name and any optional command operands, and pressing the PF key you want to associate with the report.

For example, to associate PF key 10 with the device report for devices with numbers 150 through 350 and device 370 you specify:

```
#DEV NUMBER(150:350,370)
```

and press PF key 10.

To request the report, press key 10. For the duration of your session, PF key 10 is associated with this report. Note that RMF displays the default operand menu after each # command so that you can confirm your PF key settings.

Set Delta Mode

To set the delta mode for the session, use the delta command. The syntax is:

$$D \begin{Bmatrix} \underline{\texttt{ON}} \\ \texttt{OFF} \end{Bmatrix}$$

When delta mode is in effect, certain fields in some reports, such as the processor (CPU) time in the address space state report, reflect values that show the change since the previous invocation of the report. The first request for the report will show the value RMF detects at that time; all subsequent invocations of the report will show only the change since the previous report.

Delta mode is set off when the session begins. You must enter D ON or D to set the session in delta mode. If you later want to turn off delta mode, enter D OFF. All report fields that can reflect either session or delta values will then reflect session totals rather than changes.

Set Hardcopy Mode

To set the hardcopy mode for the session, use the hardcopy command. The syntax is:

$$H \begin{Bmatrix} \underline{\texttt{ON}} \\ \texttt{OFF} \end{Bmatrix}$$

When hardcopy mode is in effect, all data in all reports requested during the session is written to the output data set.

A single output data set is created for all hardcopy and print command output generated during a single session. Because reports requested might be multi-frame reports and you might choose not to frame through all of the data during the display session, the data written to the output data set when hardcopy mode is in effect can be more extensive than that displayed on the screen.

Hardcopy mode is turned off when the session begins. You must enter H ON or H to set hardcopy mode on. The command takes effect with the next report you enter. Thus, if you decide in the middle of inspecting a multi-frame table report that you would like a printed copy of the report, you must either use the print command (P) for each frame or set hardcopy mode on and then recall the report.


```
                    RMF DISPLAY MENU
NAME      PFK#  DESCRIPTION
                                                    IPS = IEAIPSXI
ARD        1    ADDRESS SPACE RESOURCE DATA
ASD        2    ADDRESS SPACE STATE DATA
ASRM       3    ADDRESS SPACE SRM DATA
CHANNEL    4    CHANNEL PATH DISPLAY
DDMN       5    SYSTEM DOMAIN DISPLAY
DEV        6    SYSTEM DEVICE DATA
PGSP       7    SYSTEM PAGING SPACE DATA
SENQ       8    SYSTEM ENQUEUE CONTENTION
SENQR      9    SYSTEM ENQUEUE RESERVE
SPAG      10    PAGING DATA
SRCS      11    REAL STORAGE / CPU / SRM DATA
TRX       12    TRANSACTION ACTIVITY DATA
ARDJ            RESOURCE DATA FOR SPECIFIC JOBNAME
ASDJ            STATE DATA FOR SPECIFIC JOBNAME
ASRMJ           SRM DATA FOR SPECIFIC JOBNAME
DEVV            SYSTEM DEVICE DATA FOR A SPECIFIC VOL/NUMBER
IOQUEUE         I/O QUEUING ACTIVITY DATA
USER            USER PICTURE
```

Figure 4-2. Contents of the Menu Frame

Request a Report

To request a report, you enter the menu item name shown in the menu frame and any operands you require. The syntax is:

```
mm [operands]
```

Where **mm** is the menu item for the report and **operands** represents any operands needed for the report. At least one blank must appear between the menu item and any operands you specify. The menu items, possible operands, and a brief description of each report provided by RMF are shown in Figure 4-3. Each of the menu items provided by RMF is described in more detail under "Menu Items" later in this chapter.

If the menu item has a PF key associated with it, you need only enter any operand associated with the report and then press the appropriate PF key. If there are no operands associated with the report, or you wish to use the defaults, simply press the appropriate key.

Menu Item and Operands	Report Description
ASD class,status,domain	Address space state data
ASDJ jobname	Address space state data by jobname
ARD class,status,domain	Address space resource data
ARDJ jobname	Address space resource data by jobname
ASRM class,status,domain	Address space SRM services data
ASRMJ jobname	Address space SRM services data by jobname
CHANNEL	Channel path activity data
DDMN	Domain activity data
DEV type, number, or volser	Device activity data in table form
DEVV number or volser	Device activity data in row form
PGSP type	Page swap activity data
SENQ level,name	Enqueue contention activity
SENQR level	Reserve activity
SPAG	Paging
SRCS	Real storage/processor/SRM
TRX subsys, perf grps	Transaction activity data
IOQUEUE type or number	I/O Queuing Activity
USER	Reserved for an installation-supplied report

Figure 4-3. Menu Items

RMF generates two types of reports: table reports and row reports.

Table Reports: Table reports consist of up to two lines of header information and a variable number of data lines. A maximum of 21, 29, or 40 lines (depending on the terminal model being used) can appear on the screen. When a table report has more than the maximum number of lines for your device, you can use the frame command (F) to frame through all the data in the report.

Row Reports: Row reports consist of up to two lines of header information and one line of data. When you request a row report repetitively with no intervening requests for another report, the header lines and one data line appear for the first request. Each subsequent request causes an additional data line to be displayed. The newest data line is highlighted. Repetitive requests for the same report can thus be used to build a table of information. When the screen is full, the next request will cause the display to wrap around. The new line overlays the first data line of the existing display, and the new line is highlighted. The process is repeated for the second and all subsequent lines that wrap around. Any lines that are overlaid will be lost unless the report is printed.

Repeating Reports

To request reports repeatedly, use the timed update command. The syntax is:

```
T x,y
```

The timed command allows you to define the number of times you want to have a report repeated (x) and the number of seconds between reports (y). If x or y is not specified, the default for x is 10 and the default for y is 4; that is, the report is repeated 10 times at 4-second intervals. The maximum value that can be specified for x or y is 99.

To use the timed update command, you enter the command in the command input area. After you enter the command, RMF prompts you to enter the menu item of the report to be repeated. To terminate the repeated requests for a report, press the attention key. However, when you use the timed update command during a local 3270 display session, choose your options carefully; you cannot terminate the repeated reports for a local 3270 display session.

Frame Through a Table Report

To display all of the data contained in a multi-frame table report (a report that consists of more than the maximum number of lines for your device), use the frame command. The syntax is:

```
F
```

When RMF displays the first frame of a multi-frame report, a frame command (F) automatically appears in the input area. After inspecting the data in the current frame, press the enter key to see the next frame. Continue the process until you have seen all of the data that you require. If you decide at any point that you do not need to see all of the frames in a report, blank out the frame command or enter a new command and press the enter key. When RMF displays the last frame in the report, the input area is blank.

Figure 4-4 shows an example of screen output for a multi-frame report. Note that the input area is blank when the last frame is displayed.

If you enter F when there are no subsequent frames, the command is ignored.

```
F        MIG=***          CPU= 88 UIC=255 PDT=  26 DPR=  0     CHANNEL  TH
  13:08:09   CHANNEL    %CHP        CHANNEL    %CHP        CHANNEL    %CHP
             PATH TYPE  BUSY        PATH TYPE  BUSY        PATH TYPE  BUSY
              00   BY   0.00         01   BL   0.02         02  BL   0.44
              03   BL   9.94         04   BL   0.00         05  BL   0.13
              06   BL   0.00         07   BL   0.00         08  BL   0.00
              09   BL  20.76         0A   BL   0.00         0B  BL   0.02
              0C   BL   0.00         0D   BL   0.01         0E  BL   0.00
              0F   BL   0.28         10   BL   0.00         11  BL  14.06
              12   BL   0.00         13   BL   0.00         14  BL   9.23
              15   BL   0.00         16   BL   0.34         17  BL   9.06
              18   BL   0.88         19   BL   7.93         1A  BL  13.35
              1B   BL   7.93         1C   BL  28.62         1D  BL   5.59
              1E   BL   8.89         1F   BL   0.03         20  BL   0.00
              21   BL   1.49         22   BL   0.00         23  BL   0.00
              24   BL  12.48         25   BL   0.00         26  BL   0.00
              27   BL   0.00         28   BL  10.81         29  BL   0.01
              2A   BL   0.00         2B   BL   1.02         2C  BL   1.23
              2D   BL   0.00         2E   BL   9.00         2F  BL  20.48
              40   BL   0.00         41   BL   0.00         42  BL   0.32
              43   BL   5.78         44   BL   0.00         45  BL   0.00
              46   BL   0.00         47   BL   0.00         48  BL   0.00
              49   BL  16.65         4A   BL   0.00         4B  BL   0.00
              4C   BL   0.00         4D   BL   0.00         4E  BL   0.00
```

```
         MIG=***          CPU= 88 UIC=255 PDT=  26 DPR=  0     CHANNEL  TH
  13:08:09     CHANNEL    %CHP        CHANNEL    %CHP        CHANNEL    %CHP
               PATH TYPE  BUSY        PATH TYPE  BUSY        PATH TYPE  BUSY
                4F   BL   0.18         50   BL   0.00         51  BL  13.41
                52   BL   0.00         53   BL   0.00         54  BL   2.88
                55   BL   0.00         56   BL   1.10         57  BL   5.68
                58   BL   1.62         59   BL  17.13         5A  BL  10.14
                5B   BL  21.66         5C   BL   7.43         5D  BL   5.91
                5E   BL   7.15         5F   BL   0.03         60  BL   0.00
                61   BL   0.22         62   BL   0.00         63  BL   0.00
                64   BL   2.90         65   BL  10.98         66  BL   0.00
                67   BL   0.00         68   BL  11.24         69  BL   0.03
                6A   BL  23.30         6B   BL   0.91         6C  BL   2.55
                6D   BL   0.00         6E   BL  14.85         6F  BL   0.00
```

Figure 4-4. Use of the Frame Command

Recall the Previous Data

To cause the most recently displayed set of data (either a full table for a table report or a single line for a row report) to be displayed again, use the recall command. The syntax is:

```
Rmm [operands]
```

where **mm** is the menu item name and **operands** can be either the operands you specified when you requested the report or a subset of these operands. For example, if you entered ASD A,I,A, you could enter RASD T to see only the TSO address spaces or RASD B to see only the batch address spaces. No blanks can intervene between R and the menu item name. At least one blank must appear between the menu item name and any operands you require.

Example of Recalling a Report

If you requested the address space state report for a job named JOB1 with the command:

```
ASDJ JOB1
```

You would recall the report with the command:

```
RASDJ JOB1
```

Print a Report

To print the data currently displayed on the screen, use the print command.

The syntax is:

```
P
```

Because the command causes only the current screen image to be printed, you would have to use repetitive frame commands and print commands to print all of the data in a multi-frame table report.

Use this command when you want to print a single report and the session is not in hardcopy mode. In hardcopy mode, the entire report is automatically printed, and you would not need to use the print command.

A single output data set is created for all print command and hardcopy output generated during a single session.

Stop the Session

To stop the session, use the stop command. The syntax is:

```
Z
```

When you enter Z, RMF ends the session.

Menu Items

This section describes, in alphabetical order, the menu items provided by RMF. When you enter one of these menu items during a display session, the corresponding report is displayed on the screen. If your installation adds reports to those provided by RMF, you might want to note the menu item, any operands and their meanings, and the report description in this section.

ARD [class, status,domain]

Requests the report on address space resource data; the **class, status**, and **domain** operand fields specify the selection criteria for the address spaces to be included in the report.

The meaning of each operand field is:

class

A	specifies all address spaces
B	specifies batch, started task, and mount task address spaces
T	specifies TSO address spaces

status

A	specifies all address spaces
I	specifies address spaces that are active; that is, those address spaces that are currently executing, non-swappable, or swapped out and eligible for swap-in

domain

A	specifies all address spaces
nnn	specifies those address spaces in the domain identified by **nnn.**

The **class, status**, and **domain** operand fields are optional; when all are omitted, the report includes all address spaces in the system that are currently active.

The operand fields are positional; thus, any operand omitted, except the last to be specified, must be replaced by a comma; RMF uses the default value for any omitted operand.

Defaults: The menu defaults are specified in the RMF control section (ERBFMENU) that contains the menu list for display reports. Separate defaults are set for the data gathering routine and the data reporting routine. For the data gathering routine, the value for **class** is A, and the value for **status** is A. For the data reporting routine, the value for **class** is A, and the value for **status** is I. Because data is gathered for all address spaces, you can recall data for either batch, started task and mount task address spaces, or TSO address spaces without re-executing the report. You can change the menu defaults to fit the needs of your particular installation as described in Chapter 7 under "Installing a User Report" and "Using the PICTURE Macro."

Default values are also included in the data gathering routine and the data reporting routine. The "hard-coded" default for **class** is A, for **status** is A, and for **domain** is A.

To determine the operands for the report, RMF first uses the operands that you specify on the report request. If you do not specify all three operands, RMF uses the menu defaults. If the operands are still not complete, RMF uses the default values coded in the data gathering and data reporting routines.

Examples of Using Menu Items

1. To report on all address spaces that are currently active in domain 3, enter

   ```
   ARD ,,3
   ```

2. To report on all TSO address spaces that are currently active in any domain, enter

   ```
   ARD T
   ```

3. To report on all batch, started task, and mount task address spaces that are currently active in domain 3, enter

   ```
   ARD B,,3
   ```

Menu Item Commands

ARDJ jobname
Requests the report on address space resource data for the job specified in **jobname**.

ASD [class,status,domain]
Requests the report on address space state data; **class, status,** and **domain** specify the selection criteria for the address spaces to be included. The fields are optional; their meaning and use are as described under ARD.

ASDJ jobname
Requests the report on address space state data for the job specified in **jobname**.

ASRM [class,status,domain]
Requests the report on address space SRM data; **class, status,** and **domain** specify the selection criteria for the address spaces to be included. The fields are optional; their meaning and use are as described under ARD.

ASRMJ jobname
Requests the report on address space SRM data for the job specified in **jobname**.

CHANNEL
Requests the report on channel path activity data. To obtain a complete report, RMF must be initialized.

DDMN
Requests the report on domain activity data.

DEV [type]
Requests a table report on I/O device activity, where **type** can be any one of the following:

- A device class:

<u>DASD</u>	Direct access storage devices
TAPE	Magnetic tape devices
COMM	Communication equipment
CHRDR	Character reader devices
UNITR	Unit record devices
GRAPH	Graphic devices

In contrast with the Monitor I session I/O device activity report, when you specify another device class, volume serial number, or device number, you do not get DASD by default.

- One or more volume serial numbers:

```
         ( aaaaaa                 )
{VOLSER} ( aaaaaa,bbbbbb,...      )
{V     } ( aaaaaa,bbbbbb:zzzzzz   )
```

  VOLSER requests specific DASD or tape devices, where aaaaaa, bbbbbb, and zzzzzz each represent a volume serial number. You can specify any combination of one volume, a list of volumes, or a range of volumes as long as your entry does not exceed 32 characters including commas and colons. When you specify a range, use a colon as a separator to indicate that the report is to consist of all volumes from aaaaaa up to and including zzzzzz.

- One or more device numbers:

```
         ( aaa           )
{NUMBER} ( aaa,bbb,...   )
{N     } ( aaa,bbb:zzz   )
```

  NUMBER requests specific device numbers, where aaa, bbb, and zzz each represent hexadecimal device numbers. Leading zeroes can be omitted. You can specify any combination of one number, a list of numbers, or a range of numbers as long as your entry does not exceed 32 characters, including commas and colons. When you specify a range of numbers, use a colon as a separator to indicate that the report is to include all of the numbers from bbb up to and including zzz.

- One or more storage group names:

```
     ( aaaaaaaa                    )
{SG} ( aaaaaaaa,bbbbbbbb,...       )
{S } ( aaaaaaaa,bbbbbbbb:zzzzzzzz  )
```

  SG requests specific storage group names, where aaaaaaaa, bbbbbbbb, and zzzzzzzz each represent 1 to 8 character names, found in SMF type 74 and 79 records for each DASD device managed by the storage management subsystem (SMS). You can specify any combination of one storage group name, a list of names, or a range of names as long as your entry does not exceed 32 characters, including commas and colons. When you specify a range of storage group names, use a colon as a separator to indicate that the report is to include all of the names from bbbbbbbb up to and including zzzzzzzz. RMF reports the devices in sequence by device number within the storage groups.

When you request this report, a Monitor I session must be measuring any device you request.

Examples:

1. To request an I/O device activity report for all magnetic tape devices, specify:

   `DEV TAPE`

2. To request an I/O device report for volumes P50002, P50003, P50004, and P50007, specify:

   `DEV V(P50002:P50004,P50007)`

3. To request a direct access device activity report for the storage groups MANF13, MANF14, MANF15, MANF16, MANF17, MANF18 MANF19, and MANF20, specify:

   ```
   DEV SG(MANF13:MANF20)
   ```

4. To request all storage groups, specify:

   ```
   DEV SG
   ```

DEVV id

In contrast to the DEV option, which allows you to report on more than one device, this option requests a row report on a specific direct access device, where **id** can be either:

{VOLSER | V} (xxxxxx)

Requests I/O device activity for the specific volume identified by volume serial number xxxxxx.

{NUMBER | N} (yyy)

Requests I/O device activity for the specific device identified by number yyy.

When you request this report, a Monitor I session must be measuring the device you request. A storage group name is reported for any DASD device that is assigned to one.

Examples:

1. To request an I/O device activity report for tape device number 580, specify:

   ```
   DEVV N(580)
   ```

2. To request an I/O device activity report for direct access device DASD01, specify:

   ```
   DEVV V(DASD01)
   ```

IOQUEUE [type]

Requests the report on I/O queuing activity grouped by logical control unit, where **type** can be any one of the following:

- A device class:

DASD	Direct access storage devices
TAPE	Magnetic tape devices
COMM	Communication equipment
CHRDR	Character reader devices
UNITR	Unit record devices
GRAPH	Graphic devices

  In contrast with the Monitor I session I/O device activity report, when you specify another device class or N, you do not get DASD by default.

- One or more logical control unit (LCU) numbers

```
{NUMBER}     ({aaa         })
{N     }     ({aaa,bbb,... })
             ({aaa,bbb:    })
```

  NUMBER requests specific logical control unit numbers, where aaa, bbb, and zzz each represent hexadecimal LCU numbers. Leading zeroes can be omitted. You can specify any combination of one number, a list of numbers, or a range of numbers as long as your entry does not exceed 32 characters, including commas and colons. When you specify a range of numbers, use a colon as a separator to indicate that the report is to include all of the numbers from bbb up to and including zzz.

When you request this report, a Monitor I session must be measuring I/O queuing activity.

Examples:

1. To request an I/O queuing activity report for LCUs representing all magnetic tape devices, specify:

   ```
   IOQUEUE  TAPE
   ```

2. To request an I/O queuing activity report for LCU numbers D, E, F, 4E, and 55, specify:

   ```
   IOQUEUE N(D:F,4E,55)
   ```

PGSP [**PAGE** (default) | **SWAP**]

Requests the report on page or swap data set activity data, where:

- PAGE - indicates page data set activity
- SWAP - indicates swap data set activity

A Monitor I session monitoring page/swap activity must be active when you request this report.

SENQ [**S** (default) | **D** | **A,sysname** | **E,sysname** | **majorname[,minorname]**]

Requests the report on enqueue contention activity. The operands describe the type of data you require. Only one operand can be specified. The meaning of each operand field is:

S

Specifies a summary report. For each resource that had contention activity, the summary report includes the scope of the resource, the number of tasks that own the resource, the number of tasks waiting for exclusive use of the resource, and the number of tasks waiting for shared use of the resource. If you specify SENQ without operands, S is the default.

D

Specifies a detail report. For each resource that had enqueue contention, the detail report includes the scope of the resource, the name and address space identifier of each job owning or waiting for the resource, and the type and status of each job's request for the resource.

A,sysname

Specifies a report that includes all resources that a specific system holds in a global resource serialization complex, where *sysname* indicates the system for which the report is requested.

This report is useful when attempting to recover an inactive processor in a global resource serialization complex. You can request this report from an active processor in the complex and determine from the report the resources that the inactive processor held.

E,sysname
Specifies a report that includes all exclusively owned resources that a specific system held in a global resource serialization complex, where *sysname* indicates the system for which the report is requested. This report is useful when attempting to recover an inactive processor in a global resource serialization complex. You can request this report from an active processor in the complex and determine from the report the resources that the inactive processor held.

majorname[,minorname]
Specifies a detail report for a specific resource that had contention. The **majorname** field, which corresponds to the **qname** field in the ENQ and DEQ macro instructions, contains the one to eight character major name of a serially-reusable resource. Optionally, the major name is followed by a comma and a minor name. The **minorname** field, which corresponds to the **rname** field in the ENQ and DEQ macro instructions, contains the minor name of the resource.

The maximum length of the field is 32 characters, including the comma. Because the major name is 1 to 8 characters, the minor name can be from 1 to 30 characters, depending on the length of the major name. If you want a report on a minor name, but the **majorname,minorname** operand exceeds 32 characters, you must specify only the major name. RMF then collects data for all resources grouped under the major name.

Note: RMF treats the single character A, D, E, or S as a request for a report. Therefore, A or E cannot be used as a major name; S or D cannot be used as a major name unless a minor name is specified also.

Examples:

1. To obtain a summary report for all resources that have contention, enter:

   `SENQ S or SENQ`

2. To obtain a detail report for all resources that have contention, enter:

   `SENQ D`

3. To obtain a report for all resources that system C303 hold in a global resource serialization complex, enter:

   `SENQ A,C303`

4. To obtain a detail report for all resources grouped under the major name of SYSCTLG, enter:

   `SENQ SYSCTLG`

5. To obtain a detail report for all resources with the major name of SYSI and the minor name of OPENUADS, enter:

   `SENQ SYSI,OPENUADS`

Recalling the Report: If you request a report on all resources that had contention activity, you can recall the data as a summary report or a detail report on all resources or as a detail report for a specific resource. However, if your original report request was for a specific resource or for a resource held by a specific system, RMF collects no data on other resources. Thus, you cannot recall the data as a summary or detail report for all resources. Instead, the recalled report contains the data that is available. For example, if the report request is SENQ SYSDW and the recall request is SENQ S, recalling the report causes RMF to format a summary report for SYSDW only. Also, if the original report request was for a specific resource and the recall request is for a different specific resource, RMF issues a message to tell you that no data is currently available to meet your selection criteria.

SENQR [**ALLVSER** | **volser**]

Requests the report on reserve activity. The operands describe the type of data you require. Only one operand can be specified. The meaning of each operand field is:

ALLVSER
: Specifies a report that describes all reserve requests. If you specify SENQR operands, ALLVSER is the default.

volser
: Specifies a report that describes the reserve requests for a particular device, where **volser** is the one to six character volume serial number of the volume mounted on the device.

Examples:

1. To obtain a report on reserve activity for all devices, enter:

   ```
   SENQR ALLVSER or SENQR
   ```

2. To obtain a reserve activity report for the device on which the volume TSO200 is mounted, enter:

   ```
   SENQR TSO200
   ```

Recalling the Report: If you request a reserve activity report for all devices, you can recall the data as a report on all resources or as a report for a specific device. However, if your original report request was for a specific device, RMF collects no data on other devices. Thus, you cannot recall the data as a report on all devices. For example, if the report request is SENQR TSO200 and the recall request is RSENQR, the recalled report contains data for TSO200 only. Also, if the original report request was for a specific device and the recall request is for a different specific device, RMF issues a message to tell you that no data is currently available to meet your selection criteria.

SPAG
: Requests the report on system paging.

SRCS
: Requests the report on real storage/processor/SRM activity.

TRX [sname,nnn[,nnn,nnn:nnn],ALLPGN]
: Requests the transaction activity report. Because the report requires data that the SRM collects only when RMF is monitoring workload activity, the report can be obtained only when a Monitor I session is active with the WKLD option in effect. Otherwise, RMF issues a message to inform you that the report is unavailable because a Monitor I session with the WKLD option is not active.

The meaning of each operand field is:

sname
Requests data for all performance groups associated with the subsystem identified by **sname**, where **sname** is a 1-4 character subsystem name that is defined in the installation control specification.

nnn[,nnn,nnn:nnn]
Requests data for one or more performance groups, where **nnn** is a 1-3 digit performance group number. Performance groups can be specified singly (nnn) or as a range (nnn:nnn), where the lower and upper bounds are separated by a colon. Each number or range of numbers must be separated by a comma.

ALLPGN
Requests data for all performance groups. ALLPGN is the default.

The operands are not positional; thus, an omitted operand does not require a comma in its place. More than one operand can be specified. While there is no reason to combine ALLPGN with any other operand, you can specify on the same request both one or more subsystem names and individual performance groups or ranges of performance groups. In this case, if you specify individually a performance group that is also associated with a subsystem, data for that performance group appears twice in the report.

The menu default for the maximum number of performance groups that can be reported in response to any one request is 504. You can change the default as described in Chapter 7 under "Installing a User Report" and "Using the PICTURE Macro" by changing the MAXRBS operand of the PICTURE macro.

Examples:

1. To request data for performance group 1 and performance groups 5 through 10, enter

   ```
   TRX 1,5:10
   ```

2. To request data for all performance groups associated with the subsystem name of ABC, which is defined in the installation control specification, enter

   ```
   TRX ABC
   ```

3. To request data for all performance groups associated with the subsystem name of ABC and performance group 50, enter

   ```
   TRX ABC,50
   ```

   Note: If performance group 50 is also associated with the subsystem name of ABC, the data for performance group 50 will appear twice in the report.

4. To request data for all performance groups, enter

   ```
   TRX ALLPGN
     or
    TRX
   ```

USER
Requests a user-specified report. Your installation must code a corresponding data-gathering module and data-reporting module before USER can take effect.

Monitor II Background Session Control

The measurement activities of a Monitor II background session are started from the operator console with a START session command. See "Starting a Monitor II Background Session."

The options that you select control session processing. Each option is described under "Monitor II Background Session Options." The options for a session can be specified in:

- The **parm** field of the START command that is issued to start the session, or
- The RMF partitioned data set member(s), normally included in SYS1.PARMLIB

When an option is not specified in either of these sources, a program default is used. The input from each of these sources is merged to form a final set of options for the session. This process is described under "Input Merge for a Background Monitor II Session."

The options selected when the session is started can be modified during the session, as described under "Modifying the Monitor II Background Session Options."

Methods of ending the session are described under "Ending a Monitor II Background Session."

Monitor II Background Session Options

This section describes, in alphabetical order, the options that can be specified for a Monitor II background session. If your installation adds reports to those provided by RMF, you might want to note here the option, any operands and their meaning, and a report description. The program defaults are underscored where appropriate.

ARD [(class,status,domain)]/<u>NOARD</u>

Specifies whether or not address space resource data is to be collected; the **class, status**, and **domain** operand fields specify the selection criteria for the address spaces to be included. The meaning of each operand field is:

class

<u>A</u>	specifies all address spaces
B	specifies batch, started task, and mount task address spaces
T	specifies TSO address spaces

status

A	specifies all address spaces
<u>I</u>	specifies address spaces that are active; that is, those address spaces that are currently executing, non-swappable, or swapped out and eligible for swap-in

domain

<u>A</u>	specifies all address spaces
nnn	specifies those address spaces in the domain identified by **nnn**.

The **class, status,** and **domain** operand fields are optional; when all are omitted, the report includes all currently active address spaces in the system. The operand fields are positional; thus, any operand omitted, except the last to be specified, must be replaced by a comma; RMF uses the default value for any omitted operand.

Defaults: The option list defaults are specified in the RMF control section (ERBBMENU) that contains the option list for the Monitor II background session measurement activities. Separate defaults are set for the data gathering routine and the data reporting routine. For the data gathering routine, the value for **class** is A, and the value for **status** is A. For the data reporting routine, the value for **class** is A, and the value for **status** is I. You can change the menu defaults to fit the needs of your particular installation as described in Chapter 7 under "Installing a User Report" and "Using the PICTURE Macro."

Default values are also included in the data gathering routine and the data reporting routine. The "hard-coded" default for **class** is A, for **status** is A, and for **domain** is A.

To determine the operands for the option, RMF first uses the operands that you specify on the option list. If you do not specify an operand, RMF uses the option list defaults. If the operands are still not complete, RMF uses the default values coded in the data gathering and data reporting routines.

Examples:

1. To select for measurement all address spaces that are currently active in domain 3, specify

   `ARD(,,3)`

2. To select for measurement all TSO address spaces that are currently active in any domain, specify

   `ARD(T)`

3. To select for measurement all batch, started task, and mount task address spaces that are currently active in domain 3, specify

   `ARD(B,,3)`

ARDJ (jobname)/NOARDJ
Specifies whether or not address space resource data for a particular job is to be collected. The job to be measured is identified in the **jobname** field.

ASD [(class,status,domain)]/NOASD
Specifies whether or not address space state data is to be collected; **class, status,** and **domain** specify the selection criteria for the address spaces to be included. The fields are used as described under ARD.

ASDJ(jobname)/NOASDJ
Specifies whether or not address space state data is to be collected for a particular job. The job to be measured is identified in the **jobname** field.

ASRM[(class,status,domain)]/NOASRM
Specifies whether or not address space SRM data is to be collected; **class, status,** and **domain** specify the selection criteria for the address spaces to be included. The fields are used as described under ARD.

ASRMJ(jobname)/NOASRMJ
Specifies whether or not address space SRM data is to be collected for a particular job. The job to be measured is identified in the **jobname** field.

CHANNEL /NOCHANNEL
Specifies whether or not channel path activity data is to be collected.

DDMN/NODDMN
Specifies whether or not domain activity data is to be collected.

DELTA/NODELTA
Specifies whether data in some reporting fields is to reflect either total values (NODELTA) or, after the first measurement, values that reflect changes since the previous measurement (DELTA). The fields that can be affected by the choice of DELTA or NODELTA are noted in the description of each report later in this chapter.

DEV [(type)] /NODEV
Specifies whether or not data for a table report on I/O device activity is to be collected, where **type** can be any one of the following:

- A device class:

DASD	Direct access storage devices
TAPE	Magnetic tape devices
COMM	Communication equipment
CHRDR	Character reader devices
UNITR	Unit record devices
GRAPH	Graphic devices

- One or more volume serial numbers:

```
          ( aaaaaa                 )
{VOLSER}  ( aaaaaa,bbbbbb,...      )
{V     }  ( aaaaaa,bbbbbb:zzzzzz   )
```

VOLSER requests specific DASD or tape devices, where aaaaaa, bbbbbb, and zzzzzz each represent a volume serial number. You can specify any combination of one volume, a list of volumes, or a range of volumes as long as your entry does not exceed 32 characters, including commas and colons. When you specify a range, use a colon as a separator to indicate that the report is to include of all volumes from aaaaaa up to and including zzzzzz.

- One or more device numbers:

```
          ( aaa            )
{NUMBER}  ( aaa,bbb,...    )
{N     }  ( aaa,bbb:zzz    )
```

NUMBER requests specific device numbers, where aaa, bbb, and zzz each represent hexadecimal device numbers. Leading zeroes can be omitted. You can specify any combination of one number, a list of numbers, or a range of numbers as long as your entry does not exceed 32 characters, including commas and colons. When you specify a range of numbers, use a colon as a separator to indicate that the report is to consist of all numbers from bbb up to and including zzz.

- One or more storage group names:

```
          ( aaaaaaaa                     )
{SG}      ( aaaaaaaa,bbbbbbbb,...        )
{S }      ( aaaaaaaa,bbbbbbbb:zzzzzzzz   )
```

SG requests specific storage group names, where aaaaaaaa, bbbbbbbb, and zzzzzzzz each represent 1 to 8 character names, found in SMF type 74 and 79 records for each DASD device managed by the Storage Management Subsystem (SMS). You can specify any combination of one storage group name, a list of names, or a range of names as long as your entry does not exceed 32 characters, including commas and colons. When you specify a range of storage group names, use a colon as a separator to indicate that the report is to include all of the names from bbbbbbbb up to and including zzzzzzzz. RMF reports the devices in sequence by device number within the storage groups.

When you specify the DEV option, a Monitor I session must be measuring any device you request.

Examples:

1. To request an I/O device activity report for all tape devices, specify:

   `DEV(TAPE)`

2. To request an I/O device activity report for tape device numbers 580 through 584, specify:

   `DEV(NUM(580:584))`

3. To request an I/O device activity report for storage group MANF01, and storage groups MANF05 through MANF15, specify:

   `DEV(SG(MANF01,MANF05:MANF15))`

DEVV (id) /NODEVV

In contrast to the DEV option, which allows you to report on more than one device, this option specifies whether or not data is to be collected for a row report on a specific direct access device, where **id** can be either:

{VOLSER | V} (xxxxxx)

Requests I/O device activity for the specific volume identified by the volume serial number xxxxxx.

{NUMBER | N} (yyy)

Requests I/O device activity for the specific device identified by the number yyy.

When you specify the DEVV option, a Monitor I session must be measuring the device you request. A storage group name is reported for any device that is assigned to one.

Examples:

1. To request an I/O device activity report for tape device number 580, specify:

   `DEVV(N(580))`

2. To request an I/O device activity report for tape TAPE01, specify:

   `DEVV(VOLSER(TAPE01))`

IOQUEUE/ (type) /NOIOQUEUE

Specifies whether or not I/O request queuing data is to be collected, where **type** can be any one of the following:

- A device class:

DASD	Direct access storage devices
TAPE	Magnetic tape devices
COMM	Communication equipment
CHRDR	Character reader devices
UNITR	Unit record devices
GRAPH	Graphic devices

- One or more logical control unit (LCU) numbers:

```
         ( aaa          )
{NUMBER} ( aaa,bbb,...  )
{N     } ( aaa,bbb:zzz  )
```

  NUMBER requests specific logical control unit numbers, where aaa, bbb, and zzz each represent hexadecimal device numbers. Leading zeroes can be omitted. You can specify any combination of one number, a list of numbers, or a range of numbers as long as your entry does not exceed 32 characters, including commas and colons. When you specify a range of numbers, use a colon as a separator to indicate that the report is to consist of all numbers from bbb up to and including zzz.

When you specify the IOQUEUE option, a Monitor I session must be measuring any I/O queuing activity.

Examples:

1. To request an I/O queuing activity report for LCU numbers representing only the magnetic tape devices, specify:

   `IOQUEUE(TAPE)`

2. To request an I/O queuing activity report for LCU numbers D, E, F, 4E, and 55, specify:

   `IOQUEUE(N(D:F,4E,55)`

MEMBER (list)

Specifies the member or members of the RMF partitioned data set that contain the options to be used for the session. The list can contain from one to five entries, separated by commas, and each entry consists of a two-character alphameric value to be appended to ERBRMF to form the member name. When more than one member is specified, input merge takes the options from the members in left to right priority order.

The default is 01, indicating member ERBRMF01 in the partitioned data set named on the IEFRDER DD statement in the RMF cataloged procedure (normally SYS1.PARMLIB). The option must not be specified within a partitioned data set member.

For a list of the options specified in the RMF partitioned data set members supplied by IBM, see "Contents of the Parmlib Members" in Chapter 2.

Each member specified must contain options appropriate for a Monitor II background session. A member containing Monitor I session options will cause syntax errors.

OPTIONS/NOOPTIONS

Specifies whether or not a list of the options to be used is to be displayed at the operator's console at the start of the session. If the list is displayed (OPTIONS specified), the operator can respond with any desired changes, except the MEMBER option.

Figure 4-5 shows an example of the console output produced when OPTIONS in effect. Each option is followed by the input source from which RMF obtained the option. The possible sources are:

- COMMAND -- the option was specified on a START or MODIFY command
- DEFAULT -- the option was taken from the program defaults.
- CHANGED -- the option was changed by RMF. A preceding message will describe the nature of the conflict and the change RMF has made.
- MEMBER -- the option was specified in a member of the RMF partitioned data set.
- REPLY -- the option was changed by the operator in response to message ERB306I.

```
ERB103I LS : OPTIONS IN EFFECT
ERB103I LS : NOCHANNEL -- DEFAULT
ERB103I LS : NOPGSP -- DEFAULT
ERB103I LS : NODEVV -- DEFAULT
ERB103I LS : NODEV -- DEFAULT
ERB103I LS : NODDMN -- DEFAULT
ERB103I LS : NOSENQR -- DEFAULT
ERB103I LS : NOSENQ -- DEFAULT
ERB103I LS : NOTRX -- DEFAULT
ERB103I LS : NOASRMJ -- DEFAULT
ERB103I LS : NOASRM -- DEFAULT
ERB103I LS : NOARD -- DEFAULT
ERB103I LS : NOSRCS -- DEFAULT
ERB103I LS : NOSPAG -- DEFAULT
ERB103I LS : NOARDJ -- DEFAULT
ERB103I LS : NOASDJ -- DEFAULT
ERB103I LS : NOIOQUEUE -- DEFAULT
ERB103I LS : SYSOUT(A) -- MEMBER
ERB103I LS : OPTIONS -- MEMBER
ERB103I LS : REPORT(DEFER) -- MEMBER
ERB103I LS : RECORD -- MEMBER
ERB103I LS : STOP(30M) -- MEMBER
ERB103I LS : SINTV(30S) -- MEMBER
ERB103I LS : NODELTA -- MEMBER
ERB103I LS : NOUSER -- MEMBER
ERB103I LS : ASD -- MEMBER
```

Figure 4-5. OPTIONS Output for a Monitor II Background Session

Notes:

- OPTIONS can be abbreviated OPTN on the operator command
- NOOPTIONS can be abbreviated NOOPTN on the operator command
- You can avoid unnecessary console output and delay in starting the session by specifying NOOPTIONS (or NOOPTN). When RMF detects any syntax errors, OPTIONS is forced to enable the operator to examine the options and make any required changes.

PGSP [**(PAGE)** | **(SWAP)**] **/NOPGSP**

Requests that data on page or swap data set activity is to be collected, where:

- PAGE - indicates page data set activity
- SWAP - indicates swap data set activity

A Monitor I session monitoring page/swap activity must be active when you specify the PGSP option.

RECORD/NORECORD

Specifies whether or not measured data is to be written to the SMF data set. In order for RECORD to take effect, the complementary SMF enabling procedures must first be performed. These procedures are described in *System Management Facilities*.

REPORT({ **REALTIME** | **DEFER** } **)/NOREPORT**

Specifies whether or not printed interval reports of the measured data are to be produced. When reports are to be produced (REPORT specified), the REALTIME or DEFER option indicates when the reports are to be printed.

When you omit the option, the default is REPORT(DEFER). If you explicitly specify REPORT, you must specify either REALTIME or DEFER; specifying REPORT without choosing REALTIME or DEFER causes a syntax error.

REALTIME indicates that the reports are to be printed at the end of the session. When REPORT(REALTIME) is in effect, reports are also printed when:

- Session options are modified to terminate the request for a particular report, or
- Session options are modified to end the generation of all reports or to replace REPORT(REALTIME) with REPORT(DEFER).

For example, assume that the options ASRM, SPAG, and REPORT(REALTIME) are in effect for an active session. If you end the request for the system paging report by replacing SPAG with NOSPAG, any accumulated paging reports will be printed. If you change REPORT(REALTIME) to REPORT(DEFER) or NOREPORT, all accumulated reports will be printed.

DEFER indicates that the reports are to be printed after RMF terminates.

SENQ [**(S)** | **(D)** | **(A,sysname)** | **(E,sysname)** | **(majorname[,minorname])**] **/NOSENQ**

Specifies whether or not enqueue contention activity is to be measured. The operands describe the type of data you require. Only one operand can be specified. The meaning of each operand field is:

S

Specifies a summary report. For each resource that had contention activity, the summary report includes the scope of the resource, the number of tasks that own the resource, the number of tasks waiting for exclusive use of the resource, and the number of tasks waiting for shared use of the resource. If you specify SENQ without operands, S is the default for the reporting routine. For the data gathering routine, the default is D.

D
Specifies a detail report. For each resource that had enqueue contention, the detail report includes the scope of the resource, the name and address space identifier of each job owning or waiting for the resource, and the type and status of each job's request for the resource.

A,sysname
Specifies a report that includes all resources that a specific system holds in a global resource serialization complex, where *sysname* indicates the system for which the report is requested. This report is useful when attempting to recover an inactive processor in a global resource serialization complex. You can request this report from an active processor in the complex and determine from the report the resources that the inactive processor held.

E,sysname
Specifies a report that includes all exclusively-owned resources that a specific system held in a global resource serialization complex, where *sysname* indicates the system for which the report is requested. This report is useful when attempting to recover an inactive processor in a global resource serialization complex. You can request this report from an active processor in the complex and determine from the report the resources that the inactive processor held.

majorname[,minorname]
Specifies a detail report for a specific resource that had contention. The **majorname** field, which corresponds to the **qname** field in the ENQ and DEQ macro instructions, contains the one to eight character major name of a serially-reusable resource. Optionally, the major name is followed by a comma and a minor name. The **minorname** field, which corresponds to the **rname** field in the ENQ and DEQ macro instructions, contains the minor name of the resource.

The maximum length of the field is 32 characters, including the comma. Because the major name is 1 to 8 characters, the minor name can be from 1 to 30 characters, depending on the length of the major name. If you want a report on a minor name, but the **majorname,minorname** operand exceeds 32 characters, you must specify only the major name. RMF then collects data for all resources grouped under the major name.

Note: RMF treats the single character A, D, E, or S as a request for a specific type of report, such as a summary report or a detail report. Therefore, A or E cannot be used as a major name; S or D cannot be used as a major name unless a minor name is specified also.

Examples:

1. To obtain summary data for all resources that have contention, enter:

   `SENQ(S) or SENQ`

2. To obtain detail data for all resources that have contention, enter:

   `SENQ(D)`

3. To obtain a report for all resources that system C303 holds in a global resource serialization complex, enter:

   `SENQ (A,C303)`

4. To obtain detail data for all resources grouped under the major name of SYSCTLG, enter:

   `SENQ(SYSCTLG)`

5. To obtain detail data for all resources with the major name of SYSI and the minor name of OPENUADS, enter:

   `SENQ(SYSI,OPENUADS)`

Post Processing: If you specify S or D on the SENQ session option, RMF collects data about all resources that had contention activity. In this case, data is available for the post processor to format into a summary or detail report for all resources or a detail report for a specific resource. However, if you identify a resource by name or by system on the SENQ session option, RMF collects data only for the resources identified. In this case, no data on other resources is available, and the post processor cannot format a summary or detail report for all resources. Instead, the post processor formats a report that contains the data that is available. For example, if the session option is SENQ(SYSDW), and the post processor option is SENQ(D), the post processor formats a detail report for SYSDW only. Also, if the session option identified a specific resource by name and the post processor option identifies a different specific resource, RMF issues a message to tell you that no data is currently available to meet your selection criteria.

SENQR [**(ALLVSER)** | **(volser)**] **/NOSENQR**

Specifies whether or not reserve activity is to be measured. The operands describe the type of data you require. Only one operand can be specified. The meaning of each operand field is:

ALLVSER
: requests data describing all reserve requests. If you specify SENQR without operands, ALLVSER is the default.

volser
: requests data describing the reserve requests for a particular device, where **volser** is the one to six character volume serial number of the volume mounted on the device.

Examples:

1. To obtain reserve activity data for all devices, enter:

   `SENQR(ALLVSER) or SENQR`

2. To obtain reserve activity data for the device on which the volume TSO200 is mounted, enter:

   `SENQR(TSO200)`

Post Processing: If you identify a specific device by specifying a volume serial number on the SENQR session option, RMF collects data only for the device identified. In this case, no data on other devices is available to the post processor, and the post processor cannot produce a reserve activity report for all devices. If, for example, the session option is SENQR(TSO200) and the post processor option is SENQR(ALLVSER), the post processor report includes data only for the device on which the volume TSO200 is mounted. Also, if the session option identifies a specific device and the post processor option identifies a different specific device, RMF issues a message to inform you that no data is currently available to meet your selection criteria.

SINTV (value[S])

Specifies the number of seconds in each measurement interval. The range is from 1 to 3600 seconds. The default is 30 seconds.

When you require a small SINTV value, RMF overhead increases and excessive swapping can result; if this should occur, you could make the RMF address space non-swappable to minimize this overhead.

SPAG/NOSPAG

Specifies whether or not paging activity is to be measured.

SRCS/NOSRCS

Specifies whether or not real storage/processor/SRM activity is to be measured.

STOP(value [M | H])/NOSTOP

Specifies the desired duration for the session in minutes (M) or hours (H). The range is from one minute to one week (168 hours or 10,080 minutes). The default value is ten minutes. If you specify a value outside the range, RMF substitutes the default value. If neither M nor H is specified, M (minutes) is assumed. NOSTOP means the session can be ended only by a STOP command.

Notes:

1. The operator STOP command can end the session at any time, regardless of the value specified for this option.
2. The STOP option applies only to a particular Monitor II background session. RMF remains active until the operator issues a STOP system command.

SYSOUT(class)

Specifies the SYSOUT class to which the formatted interval reports are directed. Class A is the default. The SYSOUT option cannot be modified during session processing.

TRX [(sname,nnn[,nnn,nnn:nnn],ALLPGN)]/NOTRX

Requests the transaction activity report. Because the report requires data that the SRM collects only when RMF is monitoring workload activity, the report can be obtained only when a Monitor I session is active with the WKLD option in effect. Otherwise, RMF issues a message to inform you that the report is unavailable because a Monitor I session with the WKLD option is not active.

The meaning of each operand field is:

sname

Requests data for all performance groups associated with the subsystem identified by **sname**, where **sname** is a 1-4 character subsystem name that is defined in the installation control specification..

nnn[nnn,nnn:nnn]

Requests data for one or more performance groups, where **nnn** is a 1-3 digit performance group number. Performance groups can be specified individually (nnn) or as a range (nnn:nnn), where the lower and upper bounds are separated by a colon. Each number or range of numbers must be separated by a comma.

ALLPGN
Requests data for all performance groups. ALLPGN is the default.

The operands are not positional; thus, an omitted operand does not require a comma in its place. More than one operand can be specified. While there is no reason to combine ALLPGN with any other operand, you can specify on the same request both one or more subsystem names and individual performance groups or ranges of performance groups. In this case, if you specify individually a performance group that is also associated with a subsystem, data for that performance group appears twice in the report.

The menu default for the maximum number of performance groups that can be reported in response to any one request is 504. You can change the default as described in Chapter 7 under "Installing a User Report" and "Using the PICTURE Macro" by changing the MAXRBS operand of the PICTURE macro. Because an SMF record cannot exceed 32K bytes, the maximum number of performance group periods that can be included in an SMF record is 504. If you change the default to exceed 504, the record written to SMF will be truncated.

Examples:

1. To request data for performance group 1 and performance groups 5 through 10, enter:

   `TRX(1,5:10)`

2. To request data for all performance groups associated with the subsystem name of ABC, which is defined in the installation control specification, enter:

   `TRX(ABC)`

3. To request data for all performance groups associated with the subsystem name of ABC and performance group 50, enter:

   `TRX(ABC,50)`

   Note: If performance group 50 is also associated with the subsystem name of ABC, the data for performance group 50 will appear twice in the report.

4. To request data for all performance groups, enter:

   `TRX(ALLPGN) or TRX`

USER/NOUSER
Specifies whether or not a user-specified activity is to be measured. Your installation must supply a corresponding data-gathering module and data-reporting module before USER can take effect.

Conflicts Between Options

After the operator enters the START session command to start the session, RMF merges the input options (see "Input Merge for the Monitor II Background Session" later in this section). After the merge is performed, it is possible that conflicts between options can arise because some options cannot be used concurrently to control a single Monitor II background session. Should any conflicts occur, RMF detects the mutually-exclusive options during input merge and selects compatible values for these options; the operator is notified of the selections made. The possible conflicts are:

Conflict	Problem	RMF Resolution
NOREPORT AND NORECORD specified	No way for installation to obtain measurement data	Change NOREPORT to REPORT (DEFER)
STOP value specified is less than INTERVAL	Indicates session termination before obtaining any data	Set STOP value equal to INTERVAL value
REPORT(DEFER) and NOSTOP specified	SYSOUT becomes cluttered with unprinted reports	Change NOSTOP to STOP set equal to INTERVAL value

Input Merge for a Monitor II Background Session

Input merge is the process RMF performs to determine what options are to be established for the session. Input merge is performed at the start of the session and whenever the MODIFY session command is used to modify the options of an active Monitor II background session. Input merge takes options from various input sources and merges them to form a list of options for the session.

When the Session Is Started

At the start of the session, input merge uses a "fill in the blanks" procedure; that is, it takes options from the three possible input sources in a defined priority order. Once an option is filled in, no options from a lower-priority input source can override that option. The input sources and their priority order are:

1. The **parm** field of the START command. Any options specified in this input source have the highest priority. If the MEMBER option is specified, the options contained in the member(s) are processed after any other options specified on the START command and before RMF proceeds to the next input source.

   The MEMBER option can be used to specify up to five members of the RMF partitioned data set (normally SYS1.PARMLIB). If the MEMBER option is used to specify more than one member, RMF takes options from the members in left to right order. If, for example, MEMBER(03,07) is specified, RMF generates the member names ERBRMF03 and ERBRMF07 and takes options from ERBRMF03 first. Thus, if ERBRMF03 specifies NOASD and ERBRMF07 specifies ASD, NOASD is the option established for the session. If a specified member does not exist, RMF continues with the next specified member, if any, or proceeds to the next input source.

2. If the MEMBER option has not been specified on the START command, RMF uses a default member name, ERBRMF01, as the second highest priority input source. Options from the default member can fill in only those options not specified earlier on a higher priority input source.
3. The program defaults. The program defaults are the lowest priority input source; that is, the program defaults can fill in only those options not filled in by a higher priority input source.

During input merge, RMF might detect mutually exclusive options (such as RECORD and NORECORD) or syntax errors. When either occurs, RMF notes the condition and issues a warning message but does not terminate the merge. When mutually exclusive options are detected, RMF takes the value associated with the first specification of the option. When an invalid value is detected for SINTV or STOP, RMF substitutes a valid value, as described for each option under "Monitor II Background Session Options." When other syntax errors are detected, RMF ignores the input option in error and uses the next valid specification from a lower priority source.

Following a merge during which an error was detected — or any merge when the OPTIONS option is in effect — RMF displays the options resulting from the merge at the operator console; the system operator can change the options, if necessary, before session processing begins.

Examples

Assume that the operator issued the following START command to start a Monitor II background session:

```
MODIFY RMF.A,START AB,DELTA,SINTV(30),MEMBER(07)
```

The input merge process uses two of the three options from the START command to begin the list of session options:

```
DELTA
SINTV(30)
```

Because MEMBER (07) is specified in the START command, RMF generates the member name ERBRMF07 and locates it in SYS1.PARMLIB. Assume that ERBRMF07 contains the following options:

```
ASD          STOP(20)
SINTV(10)    SPAG
OPTIONS      SRCS
```

RMF would add all of these options except SINTV(10) to the option list. SINTV(10) would not be used because SINTV(30) was specified on the higher-priority START command. The option list for the session is now:

```
DELTA        STOP(20)
SINTV(30)    SPAG
ASD          SRCS
OPTIONS
```

To complete the option list, input merge would then proceed to the program defaults supplied by IBM. (These defaults are indicated in the discussion of each option earlier in this chapter under "Monitor II Background Session Options.") As a result of merging the defaults with the options already specified, input merge would build a complete list of session options. The final option list would be:

```
NOASRMJ -- DEFAULT              SYSOUT(A) -- DEFAULT
NOASRM -- DEFAULT               SRCS -- MEMBER
NOARDJ -- DEFAULT               SPAG -- MEMBER
NOARD  -- DEFAULT               ASD -- MEMBER
NOASDJ -- DEFAULT               STOP(20M) -- MEMBER
NOSENQ -- DEFAULT               NOSENQR -- DEFAULT
NOUSER -- DEFAULT               DELTA -- COMMAND
NOIOQUEUE -- DEFAULT            SINTV(30) -- COMMAND
REPORT(DEFER) -- DEFAULT        OPTIONS -- MEMBER
RECORD -- DEFAULT
```

When the Session Options Are Modified

When the MODIFY session command is used to modify the options established for an active session, RMF performs an input merge process. However, the process differs from the input merge used at the start of the session. Input merge as a result of a MODIFY session command starts with the list of options previously established and uses the input sources to **override** any previously established option.

The input sources and the order in which they are processed are:

1. The **options** field of the MODIFY session command. Any options specified here will override and replace any options in effect prior to the MODIFY session command.

2. The library data source. If the **options** field of the MODIFY command includes a MEMBER option, RMF will generate the member name and use any options specified in the member to **override** and replace any options in effect prior to the processing of the library data source, including any options specified in the MODIFY command **options** field.

   When more than one member is specified, RMF processes the members in left to right order. However, because the merge process that occurs in response to a MODIFY session command uses the input sources to override any previously-established options, options from the rightmost member will override any corresponding options from a previously-processed member. For example, if MEMBER(03,07) is specified on a MODIFY command, RMF generates the member names ERBRMF03 and ERBRMF07 and takes options from ERBRMF03 first. Thus, if ERBRMF03 specifies NOASD, the merge process places the option in the list of options currently established for the session. If, however, member ERBRMF07 specifies ASD, ASD overrides the currently-established NOASD. ASD is thus the option established for the session as a result of the MODIFY command.

When input merge is complete, RMF displays the options resulting from the merge at the operator console when an error is detected or when OPTIONS is in effect. The actions RMF takes in response to errors and the response the operator can make are identical to those for input merge at the start of the session.

For example, assume that the options for a currently-active session include NOASD, SPAG, and NOSTOP, and that you want to modify these options to ASD, NOSPAG, and STOP(40M).

If you issue the command:

```
MODIFY RMF.A,MODIFY AB,ASD,NOSPAG,STOP(40M)
```

the options will be modified as you want.

If, however, member ERBRMF09 includes the options:

```
ASD
NOSPAG
NOSTOP
```

and you issue the command:

```
MODIFY RMF.A,MODIFY AB,STOP(40M),MEMBER(09)
```

RMF will:

1. Merge the input option from the command and replace NOSTOP in the current option list with STOP(40M).
2. Merge the options from ERBRMF09 with the current options list, replacing NOASD with ASD, SPAG with NOSPAG, and STOP(40M) with NOSTOP.

Thus, any options in a member will override both any current options **and** any options specified on the MODIFY session command.

Starting a Monitor II Background Session

A Monitor II background session is started by a START session command, which is passed to RMF through the input field of the system MODIFY command. The required syntax is:

```
{MODIFY} {[RMF.]id,} {START}  session-id[,parm]
{  F   } { jobname,} {  S  }
```

id
The identifier assigned to RMF when it was initialized.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

session-id
A two-character alphanumeric identifier for the session to be started. This session identifier can be any two characters except ZZ (ZZ always indicates the Monitor I session). The session identifier appears in all RMF messages pertaining specifically to the session being started. The identifier specified must be unique; it cannot be the same as an identifier already assigned to a session that is currently active.

parm
The options for the session. Each option has the form:

```
option[(value)]
```

and must be separated by a comma. If the MEMBER option is specified, any members identified in the **value** field are used during input merge to produce the list of options for the session, though the options specified in this field of the START command have a higher priority.

Examples:

1. To start a Monitor II background session when all options are to be taken from the program defaults, issue the command:

   ```
   MODIFY RMF.A,START AB
   ```

2. To start a Monitor II background session and specify that reports be produced at the end of the session and that other options be taken from the RMF partitioned data set member ERBRMF07, issue the command:

   ```
   MODIFY RMF.A,START BB,MEMBER(07),REPORT(REALTIME)
   ```

Modifying the Monitor II Background Session Options

When required, the options in effect for an active Monitor II background session can be modified by issuing a MODIFY session command specifying the options to be changed.

Issuing a MODIFY session command forces an immediate end of interval; after interval processing, those options specified in the command are changed, a message is issued, and the session resumes processing using the new set of options. The MODIFY session command is passed to RMF through the input field of the system MODIFY command. The required syntax is:

```
{MODIFY} {[RMF.]id,} {MODIFY}  session-id[,options]
{  F   } { jobname,} {  F   }
```

id
The identifier assigned to RMF when it was initialized.

jobname
The jobname assigned to RMF when it was initialized as a batch job.

session-id
The two-character alphameric identifier of the session to be modified.

options
The session options to be changed. Each option has the form:

```
option[(value)]
```

and must be separated by a comma. All options except SYSOUT can be specified; the SYSOUT option cannot be modified during a session. Note that, if you specify the MEMBER option, any options within the member will override any options specified in the command.

Note: At the start of the session, RMF uses the values in the SINTV and STOP options to calculate the number of measurement intervals for the session. If the STOP value is 30 minutes and the SINTV value is 30 seconds, there are 60 measurement intervals during the session. When you change either of these values during the session, the duration of the session can be affected. Assume that a session has been started with the STOP and SINTV values noted above. After 15 minutes, you decide to change SINTV to 20 seconds, while STOP remains at 30 minutes. RMF re-calculates the number of intervals required by dividing the STOP value by the SINTV value. In this example, the number of intervals would be 90. The duration of the session is therefore 45 minutes.

Example

To modify the options of an active Monitor II background session to add printed output to SMF record output (NOREPORT and RECORD in effect), enter the command:

```
MODIFY RMF.A,MODIFY AB,REPORT(DEFER)
```

Ending a Monitor II Background Session

A Monitor II background session can be ended in three ways:

- Issuing a system STOP command that stops all active non-TSO RMF sessions. The syntax of the command and an example are given earlier in this chapter under "RMF Termination."
- Expiration of the time value specified in the STOP option. Use of the STOP option is described under "Monitor II Background Session Options" earlier in this section.
- Issuing a STOP session command to stop a specific active session. All other active RMF sessions continue processing. This method is described here.

The STOP session command is passed to RMF through the input field of the system MODIFY command. The required syntax is:

```
{MODIFY} { [RMF.]id,} { STOP}  session-id
{  F   } { jobname, } {  P  }
```

id
The identifier assigned to RMF when it was initialized.

jobname,
The jobname assigned to RMF when it was initialized as a batch job.

session-id
The two-character session identifier assigned to the session when it was started.

Issuing the command forces an immediate end of interval for the session; after interval processing is complete, RMF issues a message and ends the session.

Example

To stop a Monitor II background session started with a session identifier of AR, issue the command:

```
MODIFY RMF.A,STOP AR
```


Monitor II Session Reports

This section describes the Monitor II session reports in the following order:

- Address space resource data
- Address space resource data by jobname
- Address space SRM data
- Address space SRM data by jobname
- Address space state data
- Address space state data by jobname
- Channel path activity
- Domain activity
- Enqueue contention activity
- I/O device activity
- I/O queuing activity
- Page/swap data set activity
- Paging
- Real storage/processor/SRM
- Reserve activity
- Transaction activity

Sample reports are given, and the meaning of each field in each report, as well as the contents, is described, including a field definition, its use, the range of values, and the field heading. All calculated numeric values in all reports are rounded to the nearest printable value, unless otherwise noted in the individual report descriptions. A K in any field indicates that you should multiply the value by 1000. All data fields in the reports are obtained directly from the corresponding SMF record image unless otherwise indicated. Those data fields that are not obtained directly are calculated from fields in the SMF record image.

A Monitor II display session report can be displayed during a Monitor II local 3270 display session or a Monitor II TSO display session. Monitor II session reports in printed form can be obtained during or at the end of a background session, and all reports or individual reports obtained during a display session can also be printed.

Each report consists of two lines of heading information, identifying the data fields included in the report. The data lines consist of a fixed number of columns, each of which corresponds to a heading field. A report can be either a row report (such as the paging report), which consists of one data line, or a table report (such as the report on address space state data), which consists of a variable number of data lines.

This section describes each Monitor II session report RMF provides. The sample reports show the display screen contents for each report. When the reports are printed, the contents are identical to the report contents shown on the screen. However, there are differences in the format of the printed output.

For a **Monitor II display session**, RMF creates a single output data set for each session. All printed output resulting from either hardcopy mode or the print command is sent to the same output data set. Each report printed when the session is in hardcopy mode is delimited by a line of plus signs. Each report printed as a result of the print display command is delimited by a line of asterisks (*). Fields within the line of delimiters indicate the operands specified on the report request, whether the session is in delta or total mode, and the name of the report.

When there are repetitive requests for the same row report, headings appear for the first request; data lines appear for each subsequent request. A field within the line indicates the time of the report.

For a **Monitor II background session,** RMF creates a single output data set for each report requested. All interval reports for that measurement activity are written to a single output data set. Thus, if you request three measurements for a session with five reporting intervals, RMF will create three data sets and write five reports to each data set.

The printed output generated for each table report requested is preceded by a line of plus signs (+). Fields within the line of plus signs indicate the option and any operands specified for the report and whether the session is in delta mode or total mode. Each iteration of the report is separated by a line of plus signs. A field within the line of plus signs indicates the time of the report. The column headings are repeated for each iteration of the report.

The printed output generated for each row report is preceded by a single line of plus signs (+). Fields within the line of plus signs indicate the option and any operands specified for the report and whether the session is in delta or total mode. Headings appear only once. There is no delimiter between successive iterations of a row report, and a field within the report line indicates the time of the report.

Figure 4-6 shows an example of printed report output generated during a Monitor II background session.

Figure 4-7 shows an example of printed output from a **display session**. The first report was printed as a result of the print display command; note that it is a screen image of a report where repetitive requests had caused the display to wrap around. The second report was printed as a result of the user's entering the H ON or H command, putting the session in hardcopy mode and causing all reports to be printed as well as displayed.


```
                                                        SESSION NAME BB

 ++++++++++    N(274)                          +++++++ TOTAL MODE +++++++++++++++++++++++++ ASD      +++++++++++++++++++

 +++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++ 14.35.46 +++++++++++++++++++
14:35:46       P P C  R  DP   RS  ESF  ESF   TAR   WS   TX  WRK  CPU  I/O STM
JOBNAME  DMN   G P L  LS PR   F        +RS   WSS   IN   SC  RV   RV   RV  RV
*MASTER*  0    0 1 NS    FF  133   0  133     0    0    0    0   +0   +0   +0
PCAUTH    7    3 1 NS    76   33   0   33     0   11    0    0   +0   +0   +0
TRACE     7    3 1 NS    77  176   0  176     0   11    0    0   +0   +0   +0
GRS       0    0 1 NS    FF  561   0  561   33K   11    0    0   +0   +0   +0
CONSOLE   7    3 1 NS    FF  174   0  174     0   11    1    0   +0   +0   +0
ALLOCAS   7    3 1 NS    71  986   0  986     0   11    1    0   +0   +0   +0
LLA       7    3 1 NS    71  170   0  170     0   30    0  150   +0   +0   +0
RMF33     7    3 1 NS    71  159   0  159     0   34    0  150   +0   +0   +0
 +++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++ 14.35.54 +++++++++++++++++++
14:35:54       P P C  R  DP   RS  ESF  ESF   TAR   WS   TX  WRK  CPU  I/O STM
JOBNAME  DMN   G P L  LS PR   F        +RS   WSS   IN   SC  RV   RV   RV  RV
DSNDBM1   7   13 1 NS    FC 1350   0 1350     0   33    0  150   +0   +0   +0
BOYLEMMM  2    1 2 IN    78   85  48  133     0    0    4  101   +0   +0   +0
SMF       0    0 1 NS    FF   94   0   94     0   11    0  150   +0   +0   +0
DFHSM     7    3 1 NS    74  610   0  610     0  527    1  150   +0   +0   +0
VTAM      7    5 1 NS    FD  678   0  678     0   34    0    0   +0   +0   +0
SOS       7    3 1 IN    79   76 491  567     0    0 1.1K  150   +0   +0   +0
AMSAQFT   7    3 1 NS    72   54   0   54     0   33    0  150   +0   +0   +0
JES2      7    4 1 NS    FE  900   0  900     0   24    0    0   +0   +0   +0
CATALOG   0    0 1 NS    FF 1552   0 1552     0   11    0  150   +0   +0   +0
 +++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++ 14.36.02 +++++++++++++++++++
14:36:02       P P C  R  DP   RS  ESF  ESF   TAR   WS   TX  WRK  CPU  I/O STM
JOBNAME  DMN   G P L  LS PR   F        +RS   WSS   IN   SC  RV   RV   RV  RV
HUBERF    2    1 2 IN    78  498  99  597     0    0    3    0   +0   +0   +0
ZAPPERD   2    1 2 IN    72  110  11  121     0    0    1  150   +0   +0   +0
PTRACYB   2    1 2 IN    78  174  55  229     0    0    3    0   +0   +0   +0
IRLMPROC  7   12 1 NS    FC   61   0   61     0   33    0    0   +0   +0   +0
```

Figure 4-6. Format of Printed Reports from a Monitor II Background Session

```
                                      RMF MONITOR II HARDCOPY LOG
                                           SESSION NAME AA
         OS/VS2                  SYSTEM ID 168A         DATE 04/15/89
         SP2.2.3                 RPT VERSION 3.5.1      TIME 10.36.58

++++++++++   N(274)                    +++++++ TOTAL MODE ++++++++++++++++++++++++++ DEVV       ++++++++++++++++++++
SG MVSXA        DEV     ACTIV RESP IOSQ --DELAY--  PEND DISC CONN %DEV %DEV
  TIME   VOLSER NUM LCU  RATE  TIME TIME  CUB   DB  TIME TIME TIME UTIL RESV I%
10:36:58 GNRL01 274 ---  0.000   0    0              0.0  0.0  0.0 .000  0.0 54
10:37:29 GNRL01 274 ---  0.000   0    0              0.0  0.0  0.0 .000  0.0  1
10:37:59 GNRL01 274 ---  0.000   0    0              0.0  0.0  0.0 .000  0.0 52
10:38:29 GNRL01 274 ---  0.000   0    0              0.0  0.0  0.0 .000  0.0  1
```

Figure 4-7. Format of Printed Reports for a Monitor II Display Session

Address Space Resource Data Report

The address space resource data report gives information on the system resources — such as processor time, paging, and real storage — used by each address space in the system or each address space that meets the selection criteria that you specify when you request the report.

The number of data lines in the report depends on the number of address space identifiers in the system that meet your selection criteria. The report is very useful in enabling you to determine which jobs are creating performance problems, and, once the problem job is identified, an address space resource data report for that particular job can be requested, thus enabling you to focus your reporting on a known problem area.

Figure 4-8 shows a sample address space resource data report. The field headings, their meanings, and the possible contents are described in Figure 4-9.

To request an address space resource data report during a display session, specify:

```
ARD [class,status,domain]
```

```
F                                     CPU= 96 UIC=255 PDT=   0 DPR=  0    ARD       T
 09:00:16 DEV     FX  LS NLS X SRM    TCB     CPU     PIN  EXCP SWAP LPA CSA NV  VIO
 JOBNAME  CONN   BEL  QA  QA M ABS    TIME    TIME    RT   RATE RATE RT  RT  RT  RT
 *MASTER*  1938    0  60   0   .00  121.77 413.13 ---- ----- ---- --- --- --- ---
 PCAUTH   .0000    0  18   0 X .00    0.03   0.03 ---- ----- ---- --- --- --- ---
 TRACE    .0000    0  94   1 X .00    0.03   0.03 ---- ----- ---- --- --- --- ---
 GRS       2028    5  22  10 X .00   49.98 320.42 ---- ----- ---- --- --- --- ---
 CONSOLE  79.98    0  14   0 X .00   43.33  52.01 ---- ----- ---- --- --- --- ---
 ALLOCAS  .0000    0  17   0 X .00    0.07   0.07 ---- ----- ---- --- --- --- ---
 LLA      3.520    0  18   0 X .00    1.41   1.78 ---- ----- ---- --- --- --- ---
 RMF33    .6090    0  18   0   .00    5.49  15.24 ---- ----- ---- --- --- --- ---
 VTAM      3147    1  24   1   .00   67.33 104.66 ---- ----- ---- --- --- --- ---
 DSN2DBM1 2.144    0  35  14 X .00    0.53   0.68 ---- ----- ---- --- --- --- ---
 SMF      11.72    0  16   0 X .00    0.27   3.05 ---- ----- ---- --- --- --- ---
 DFHSM     2009    3  34   1   .00  330.10 358.71 ---- ----- ---- --- --- --- ---
 AMSAQXA  3.637    0  15   0   .00    3.38   3.83 ---- ----- ---- --- --- --- ---
 JES2     832.2   30  30  30   .00  399.74 453.63 ---- ----- ---- --- --- --- ---
 CATALOG  308.3    0  72   0 X .00  354.47 368.26 ---- ----- ---- --- --- --- ---
 ZAPPERF  261.5    2  32   5   .00  196.95 210.39 ---- ----- ---- --- --- --- ---
 D53NJV1B .1280    1  19   1   9.2    0.08   0.09 ---- ----- ---- --- --- --- ---
 DSNMSTR  11.10    0  28   1 X .00    7.30  12.91 ---- ----- ---- --- --- --- ---
 DSN2MSTR 1.745    0  29   1 X .00    6.72  10.97 ---- ----- ---- --- --- --- ---
 IRLMPRC2 .1480    0  14   0 X .00    0.06   0.29 ---- ----- ---- --- --- --- ---
 DSNDBM1  42.38    0  52   6 X .00   11.77  15.93 ---- ----- ---- --- --- --- ---
```

Figure 4-8. Address Space Resource Data (ARD) Report

Field Heading	Meaning
hh.mm.ss JOBNAME	One to eight character jobname associated with the address space. The JOBNAME heading is preceded by the time the report was requested, in the form hh.mm.ss
DEV CONN	Five-digit number (with a decimal point that floats to achieve maximum precision) indicating, in seconds, the device connect time used by the job. If the device connect time is greater than 99999 seconds, then it is expressed in hours (a four-digit number with a floating decimal point followed by H). If the device connect time exceeds 76.4 hours, three asterisks will appear in the field. Depending on your choice, this field can contain either a total value for the address space or a changed (delta) value since the last report request.
FX BEL	One to three digit number indicating the number of fixed pages below 16-megabytes for the job.
LS QA	One to three digit number indicating the number of private LSQA fixed pages for the job.
NLS QA	One to three digit number indicating the number of private non-LSQA fixed pages for the job.
XM	Cross memory address space indicator. When the field contains XM, the line of data describes a cross memory address space; that is, an address space accessed primarily from other address spaces by means of cross memory functions.
SRM ABS	Number in the form nnn, indicating the total SRM service absorption rate for the job. This field is reported only for address spaces that are currently in main storage. Dashes are used when no data is reported. Depending on your choice, this field can contain either a total value for the RMF session or a changed (delta) value since the last report request.
TCB TIME	The number of seconds of TCB processor time used by the current job step. Depending on your choice, this field can contain either a total value for the RMF session or a changed (delta) value since the report request.
CPU TIME	Amount of processor (TCB + SRB) time, in seconds, for the current job step, in the form nnn.dd, where nnn represents the number of seconds and dd hundredths of a second. Depending on your choice, this field can contain either a total value for the current job step or a changed (delta) value since the last report request. When a valid delta value cannot be computed because the job has changed steps between requests for the report, this field contains dashes when delta mode is in effect.

PIN RT	One to four digit number indicating the page-in rate (PIN).
	
	RMF calculates the page-in rate in one of two ways:
	
	• For cross-memory address spaces, the calculation used is
	
	`PIN=  number of page-ins / transaction residency time`
	
	where the transaction residency time is equal to the (current TOD-OUCBTMS)+OUXBTRR.
	
	• For all other address spaces, the calculation used is
	
	`PIN=  number of page-ins / seconds of CPU time`
	
	This field always contains the rate since the last report request. An asterisk (*) in this field indicates that RMF is unable to calculate a value.
	
	If the transaction requires storage isolation, the value reported can be used to establish initial threshold values and to evaluate the effectiveness of these values.

Figure 4-9 (Part 1 of 2). Fields in the Address Space Resource Data (ARD) Report

Field Heading	Meaning
EXCP RATE	One to five digit number indicating the EXCP rate. This field always contains the rate since the last report request.
SWAP RATE	One to four digit number indicating the page rate (the sum of pages in and pages out) for the job. This field always contains a value that reflects the change since the last report request. It is reported only for address spaces that are currently in main storage.
LPA RT	One to three digit number indicating the common LPA page-in rate for the current transaction. This field always contains a value that reflects the change since that last report request. It is reported only for address spaces that are currently in main storage.
CSA RT	One to three digit number indicating the common CSA page-in rate for the current transaction. This field always contains a value that reflects the change since the last report request. It is reported only for address spaces that are currently in main storage.
NV RT	One to three digit number indicating the private non-VIO page rate (the sum of pages in and pages out) for the current transaction. This field always contains a value that reflects the change since the last report request. It is reported only for address spaces that are currently in main storage.
VIO RT	One to three digit number indicating the private VIO page rate (the sum of pages in and pages out) for the current transaction. This field always contains a value that reflects the change since the late report request. It is reported only for address spaces that are currently in main storage.

Figure 4-9 (Part 2 of 2). Fields in the Address Space Resource Data (ARD) Report

Address Space Resource Data by Jobname Report

The address space resource data by jobname report is the same as the address space resource data report, except that it consists of one line of data that presents address space resource data for the particular jobname that you request. The format and the information for the reports are identical, with one exception: the heading for the first field in the jobname report is

cccccccc

TIME

where **cccccccc** is the one to eight character name of the job you requested. The data in the first field of the report is the time the report was requested, in the form hh:mm:ss.

You can repeat the requests for the report to build a table of information about the address space resource data for a particular job. Figure 4-10 shows an example of repetitive requests for the report.

To request an address space resource data by jobname report during a display session, specify:

```
ARDJ jobname
```

```
                                           CPU= 39 UIC=255 PDT=   6 DPR= 23    ARDJ    T
DSNMSTR  DEV     FX  LS NLS X SRM     TCB     CPU     PIN   EXCP SWAP LPA CSA NV   VIO
  TIME   CONN   BEL  QA  QA M ABS     TIME    TIME    RT    RATE RATE RT  RT  RT   RT
08:53:20 5.619   0   28   1 X .00     2.04    3.24 ----  ----- ---- --- --- --- ---
08:53:25 5.619   0   28   1 X .00     2.04    3.24  0.0  0.00 0.00 .00 .00 .00 .00
08:53:34 5.619   0   28   1 X .00     2.04    3.25  0.0  0.00 0.00 .00 .00 .00 .00
08:53:44 5.619   0   28   1 X .00     2.05    3.26  0.0  0.00 0.00 .00 .00 .00 .00
08:53:46 5.619   0   28   1 X .00     2.05    3.26  0.0  0.00 0.00 .00 .00 .00 .00
08:53:56 5.619   0   28   1 X .00     2.07    3.28  0.0  0.00 0.00 .00 .00 .00 .00
08:53:57 5.619   0   28   1 X .00     2.07    3.28  0.0  0.00 0.00 .00 .00 .00 .00
08:53:58 5.619   0   28   1 X .00     2.07    3.29  0.0  0.00 0.00 .00 .00 .00 .00
08:53:59 5.619   0   28   1 X .00     2.07    3.29  0.0  0.00 0.00 .00 .00 .00 .00
08:53:59 5.619   0   28   1 X .00     2.07    3.29 ----  ----- ---- --- --- --- ---
08:54:00 5.619   0   28   1 X .00     2.07    3.29  0.0  0.00 0.00 .00 .00 .00 .00
08:54:02 5.619   0   28   1 X .00     2.07    3.29  0.0  0.00 0.00 .00 .00 .00 .00
08:54:08 5.619   0   28   1 X .00     2.08    3.29  0.0  0.00 0.00 .00 .00 .00 .00
08:54:09 5.619   0   28   1 X .00     2.08    3.29  0.0  0.00 0.00 .00 .00 .00 .00
08:54:56 5.619   0   28   1 X .00     2.12    3.35  0.0  0.00 0.00 .00 .00 .00 .00
08:54:58 5.619   0   28   1 X .00     2.12    3.35  0.0  0.00 0.00 .00 .00 .00 .00
```

Figure 4-10. Address Space Resource Data by Jobname (ARDJ) Report

Address Space SRM Data Report

The address space SRM services data report gives an overview of the SRM services — such as processor service, MSO service, and I/O service — used by each address space in the system or each address space that meets the selection criteria that you specify when you request the report.

The number of data lines in the report depends on the number of address space identifiers in the system that meet your criteria. The report can help you to determine which jobs are using which services and whether certain jobs are creating performance problems by making excessive use of system services.

Figure 4-11 shows a sample address space SRM services data report. The field headings, their meanings, and the possible contents are described in Figure 4-12

To request an address space SRM data report during a display session, specify:

```
ASRM [class,status,domain]
```

```
F                                  CPU= 49 UIC=255 PDT=  69 DPR= 27    ASRM     T
11:40:23        P P  TRANS    TRANS   TX   TX     TX      TX      TX      TX   SESS
JOBNAME  DMN    G P  ACTIVE   CUR RES CT   SC    CPU     MSO     IOC     SRB  TOTAL
*MASTER*  0     0 1 45:04:52 45:04:52  1    0 3.153M 2.511M 814589 15.94M 22.42M
PCAUTH    7     3 1 45:05:02 45:05:02  1    0      1     16      0      6     23
TRACE     7     3 1 45:05:02 45:05:02  1    0      1     24      0     14     39
GRS       0     0 1 45:05:02 45:05:02  1    0 2.550M 7.372M     80 10.08M 20.00M
CATALOG   0     0 1 45:05:02 45:05:02  1    0 24.14M 153.1M 1.779M 1.037M 180.1M
CONSOLE   7     3 1 45:05:01 45:05:01  1    1 1.245M 496543 354702 205799 2.302M
ALLOCAS   7     3 1 45:04:53 45:04:53  1    1    251    833     10     13   1107
LLA       7     3 1 44:56:23 44:56:23  1    0  15329  17812  18900   4232  56273
RMF33     7     3 1 44:56:23 44:56:23  1    0 1.341M 1.233M   2620 1.014M 3.591M
NETMON    7     3 1 22:26:18 22:26:18  1    1  10434  12864     95  20573  43966
VTAM      7     5 1 44:55:10 44:55:10  1    0 4.086M 18.70M  74003 4.273M 27.13M
SMF       0     0 1 45:04:09 45:04:09  1    0   1816    693    245  22105  24859
AMSAQXI   7     3 1 44:56:22 44:56:22  1    0  76284  16807  19035   5842 117968
JES2      7     4 1 45:04:04 45:04:04  1    0 15.60M 79.64M 7.903M 2.066M 105.2M
D24AJD1J  2     1 2 00:36:00 00:35:59  1    1 275054 161795 289713  63643 790205
NLDM      7     3 1 24:39:32 08:10:14  1  *** 3.085M 10.09M 873871 406436 14.46M
VOYAGER1  7     3 1 28:10:28 28:10:28  1    0  27415  43721   6235  20463  97834
SONORA11  7     3 1 24:14:14 24:14:14  1    0 336536 562881 188174  77631 1.165M
RMFGAT    9     8 1 27:15:25 27:15:25  1    0 6.015M 24.69M 196504 156837 31.06M
DSNDBM1   7    13 1 26:31:00 26:31:00  1    0 757660 7.671M  13245 647118 9.089M
DSN2DBM1  7     3 1 26:25:00 26:25:00  1    0  22757 143082   2400   3767 172006
```

Figure 4-11. Address Space SRM Data (ASRM) Report

Field Heading	Meaning
hh.mm.ss JOBNAME	One to eight character jobname associated with the address space. The JOBNAME heading is preceded by the time the report was requested, in the form hh.mm.ss.
DMN	One to three digit number identifying the domain associated with the job.
P G	One to three digit number identifying the performance group associated with the job.
P P	One digit number identifying the performance period associated with the job
TRANS ACTIVE	Transaction elapsed time; that is, the time that has elapsed since the current transaction began, in the form hh:mm:ss.
TRANS CUR RES	Time, in the form hh.mm.ss, elapsed since the current transaction became resident.
TX CT	One to three digit number indicating the transaction count for the job. This field is reported only for address spaces that are currently in storage. Dashes are used when no data is reported.
TX SC	One digit number indicating the swap count for the current transaction.
TX CPU	One to six digit number indicating the processor service consumed by the current transaction. The number can be followed by an M; the M indicates millions of units absorbed.
TX MSO	One to six digit number indicating the MSO service consumed by the current transaction. The number can be followed by an M; the M indicates millions of units absorbed.
TX IOC	One to six digit number indicating the I/O service consumed by the current transaction. The number can be followed by an M; the M indicates millions of units absorbed.
TX SRB	One to six digit number indicating the SRB service consumed by the current transaction. The number can be followed by an M; the M indicates millions of units absorbed.
SESS TOTAL	One to six digit number indicating the total SRM services consumed by the entire job. The number can be followed by an M; the M indicates millions of units absorbed.

Figure 4-12. Fields in the Address Space SRM Data (ASRM) Report

Address Space SRM Data by Jobname Report

The address space SRM data by jobname report is the same as the address space SRM data report, except that it consists of one line of data that presents address space SRM data for the particular jobname that you request. The format and the information for the reports are identical, with one exception: the heading for the first field in the jobname report is

cccccccc

TIME

where **cccccccc** is the one to eight character name of the job you requested. The data in the first field of the report is the time the report was requested, in the form hh.mm.ss.

You can repeat the requests for the report to build a table of information about the address space SRM data for a particular job. Figure 4-13 shows an example of repetitive requests for the report.

To request an address space SRM data by jobname report during a display session, specify:

```
ASRMJ jobname
```

```
                                  CPU= 62 UIC=255 PDT=  58 DPR= 22   ASRMJ  T
ALLOCAS         P P  TRANS   TRANS    TX   TX    TX     TX     TX     TX   SESS
  TIME   DMN    G P  ACTIVE  CUR RES  CT   SC    CPU    MSO    IOC    SRB  TOTAL
08:57:48   7    3 1 00:41:49 00:41:49   1    1   263    866     10      7   1146
08:57:51   7    3 1 00:41:51 00:41:51   1    1   263    866     10      7   1146
08:58:00   7    3 1 00:42:00 00:42:00   1    1   263    866     10      7   1146
08:58:23   7    3 1 00:42:24 00:42:23   1    1   263    866     10      7   1146
08:58:25   7    3 1 00:42:26 00:42:26   1    1   263    866     10      7   1146
08:58:27   7    3 1 00:42:27 00:42:27   1    1   263    866     10      7   1146
08:58:28   7    3 1 00:42:29 00:42:28   1    1   263    866     10      7   1146
08:58:29   7    3 1 00:42:30 00:42:30   1    1   263    866     10      7   1146
08:58:38   7    3 1 00:42:39 00:42:39   1    1   263    866     10      7   1146
08:58:40   7    3 1 00:42:40 00:42:40   1    1   263    866     10      7   1146
08:58:41   7    3 1 00:42:42 00:42:42   1    1   263    866     10      7   1146
08:59:38   7    3 1 00:43:38 00:43:38   1    1   263    866     10      7   1146
08:59:50   7    3 1 00:43:50 00:43:50   1    1   263    866     10      7   1146
```

Figure 4-13. Address Space SRM Data by Jobname (ASRMJ) Report

Address Space State Data Report

The address space state data report gives an overview of the current state of each address space in the system or each address space that meets the selection criteria that you specify when you request the report. Basically, it tells you where each address space is and what it is doing; you can use the report, for example, to determine which jobs are using large amounts of real storage or which jobs are being swapped excessively and why the swapping is occurring. Once a particular job has been spotted as a problem, address space state data reports for that particular job can be requested, thus enabling you to focus your reporting on a known problem area.

The number of rows in the report depends on the number of address space identifiers that meet your selection criteria. Figure 4-14 shows a sample report. The field headings, meanings, and contents are described in Figure 4-15. Information about SRM service is available in the address space SRM data (ASRM) report.

To request an address space state data report during a display session, specify:

```
ASD [class,status,domain]
```

```
                           MIG=***  CPU= 68 UIC=255 PDT=  50 DPR=  0    ASD      T
11:01:06          P P C  R  DP   RS  ESF  ESF   TAR   WS   TX  WRK  CPU  I/O STM
JOBNAME  DMN      G P L  LS PR   F         +RS  WSS   IN   SC  RV   RV   RV  RV
*MASTER*  0       0 1 NS    FF  133    0  133     0    0    0    0   +0   +0   +0
PCAUTH    7       3 1 NS    76   33    0   33     0   11    0    0   +0   +0   +0
TRACE     7       3 1 NS    77  176    0  176     0   11    0    0   +0   +0   +0
GRS       0       0 1 NS    FF  561    0  561   33K   11    0    0   +0   +0   +0
CONSOLE   7       3 1 NS    FF  174    0  174     0   11    1    0   +0   +0   +0
ALLOCAS   7       3 1 NS    71  986    0  986     0   11    1    0   +0   +0   +0
LLA       7       3 1 NS    71  170    0  170     0   30    0  150   +0   +0   +0
RMF33     7       3 1 NS    71  159    0  159     0   34    0  150   +0   +0   +0
DSNDBM1   7      13 1 NS    FC 1350    0 1350     0   33    0  150   +0   +0   +0
BOYLEMMM  2       1 2 IN    78   85   48  133     0    0    4  101   +0   +0   +0
SMF       0       0 1 NS    FF   94    0   94     0   11    0  150   +0   +0   +0
DFHSM     7       3 1 NS    74  610    0  610     0  527    1  150   +0   +0   +0
VTAM      7       5 1 NS    FD  678    0  678     0   34    0    0   +0   +0   +0
SOS       7       3 1 IN    79   76  491  567     0    0 1.1K  150   +0   +0   +0
AMSAQFT   7       3 1 NS    72   54    0   54     0   33    0  150   +0   +0   +0
JES2      7       4 1 NS    FE  900    0  900     0   24    0    0   +0   +0   +0
CATALOG   0       0 1 NS    FF 1552    0 1552     0   11    0  150   +0   +0   +0
HUBERF    2       1 2 IN    78  498   99  597     0    0    3    0   +0   +0   +0
ZAPPERD   2       1 2 IN    72  110   11  121     0    0    1  150   +0   +0   +0
PTRACYB   2       1 2 IN    78  174   55  229     0    0    3    0   +0   +0   +0
IRLMPROC  7      12 1 NS    FC   61    0   61     0   33    0    0   +0   +0   +0
```

Figure 4-14. Address Space State Data (ASD) Report

Field Heading	Meaning
hh.mm.ss JOBNAME	One to eight character jobname associated with the address space. The JOBNAME heading is preceded by the time the report was requested, in the form hh.mm.ss.
DMN	One to three digit number identifying the domain associated with the job.
P G	One to three digit number identifying the performance group associated with the job.
P P	One digit number identifying the performance period associated with the job.
C L	Two character field showing the current location of the job. The possible contents and their meanings are: IN - In storage OT - Swapped out and ready LO - Logically swapped out NS - Non-swappable WM - Wait queue/MSO WL - Wait queue/long wait WT - Wait queue/terminal wait WO - Wait queue/reasons other than WM, WL, or WT DL - TSO user delayed by SRM to meet response time objective PR - Privileged >> - In the process of being swapped out of storage << - In the process of being swapped into storage
R LS	Two character field showing the reason for the last swap out associated with the job. This field is blank when the current location of the job is NS, IN, or PR. The possible contents and their meanings are: TI - Terminal input wait TO - Terminal output wait LW - Long wait XS - Auxiliary storage shortage RS - Real storage shortage DW - Detected wait RQ - Requested swap NQ - Enqueue exchange EX - Exchange based on recommendation value US - Unilateral TS - Transition Swap

Figure 4-15 (Part 1 of 2). Fields in the Address Space State Data (ASD) Report

Field Heading	Meaning
DP PR	One to two digit number indicating the ASCB dispatch priority for the job.
RS F	Indicates the number of real storage frames assigned to the job. If a job is swapped out of real storage the number represents the number of real storage frames assigned to the job before the swap out occurred.
	ESF One to three digit number indicating the number of extended storage frames occupied by virtual storage pages associated with this address space. If no extended storage is installed, this field is blank.
+RS	ESF Indicates the total number of real storage frames and extended storage frames occupied by virtual pages associated with this address space. If no extended storage is installed, this field is blank.
TAR WSS	Indicates the target working set size for the job (in number of pages).
WS IN	Indicates the size of the working set for the job (in number of pages) at the time it was last swapped in.
TX SC	One to three digit number indicating the swap count for the current transaction.
WRK RV	One to three digit number indicating the workload manager recommendation value for the job. Note that only integer values are reported; fractional values appear as zero.
CPU RV	One to three digit number, preceded by a plus (+) or minus (-) sign, indicating the resource manager CPU recommendation value for the job. Note that only integer values are reported; fractional values appear as zero.
I/O RV	One to three digit number preceded by a plus (+) or minus (-) sign, indicating the resource manager I/O recommendation value for the job. Note that only integer values are reported; fractional values appear as zero.
STM RV	One to three digit number, preceded by a plus (+) or minus (-) sign, indicating the SRM storage load balancing recommendation value for the job.

Figure 4-15 (Part 2 of 2). Fields in the Address Space State Data (ASD) Report

Address Space State Data by Jobname Report

The address space state data by jobname report is the same as the address space state data report except that it consists of one line of data giving address space state data for the particular jobname that you request. The format and the information for the reports are identical, with one exception; the heading for the first field in the jobname report is:

cccccccc

TIME

where **cccccccc** is the one to eight character name of the job you requested. The data in the first field is the time of the report, in the form of hh:mm:ss.

You can repeat the requests for the report to build a table of information about the address space state data for a particular job. Figure 4-16 shows an example of repetitive requests for the report.

To request an address space state data by jobname report during a display session, specify:

```
ASDJ jobname
```

```
                                            CPU= 39 UIC=255 PDT=  16 DPR= 23   ASDJ   T
DSNDBM1          P P C  R  DP   RS  ESF  ESF   TAR   WS   TX  WRK  CPU  I/O STM
  TIME    DMN    G P L  LS PR   F         +RS   WSS   IN   SC  RV   RV   RV  RV
08:56:20    7  13 1 NS    FB 1322               0   32    0    0   +0   +0   +0
08:56:23    7  13 1 NS    FB 1322               0   32    0    0   +0   +0   +0
08:56:27    7  13 1 NS    FB 1322               0   32    0    0   +0   +0   +0
08:56:28    7  13 1 NS    FB 1332               0   32    0    0   +0   +0   +0
08:56:30    7  13 1 NS    FB 1332               0   32    0    0   +0   +0   +0
08:56:59    7  13 1 NS    FB 1309               0   32    0    0   +0   +0   +0
08:57:01    7  13 1 NS    FB 1309               0   32    0    0   +0   +0   +0
08:56:59    7  13 1 NS    FB 1309               0   32    0    0   +0   +0   +0
08:57:01    7  13 1 NS    FB 1309               0   32    0    0   +0   +0   +0
```

Figure 4-16. Address Space State Data by Jobname (ASDJ) Report

Channel Path Activity Report

The channel path activity report gives you information about channel path loading for all channel paths in the system. The channel path is the physical interface that connects control units and devices to the I/O subsystem. Channel path activity information, I/O device activity, and I/O request queuing information can be used to identify performance bottlenecks associated with the channel paths. Requesting the report during a display session gives you the ability to see how channel path use changes during the interval.

You can obtain the report whether or not a Monitor I session measuring channel path activity is active. However, the channel path type appears only when RMF is active. The Monitor II session channel path activity report contains data for every channel path that is online at the time you request the report. RMF obtains information about channel path use from data the hardware maintains.

Your initial request establishes the base for reporting channel path activity but does not generate a report. After processing your initial request, RMF prompts you to press the ENTER key. Pressing the ENTER key causes RMF to generate the report. The reporting interval is the period between any two consecutive Monitor II requests. For the first report, the reporting interval is the period between your initial request and the time you press the ENTER key. If you continue to request the report, the reporting interval is the time between each report request, and the data reflects the channel activity since your last request.

To request a channel path activity report during a display session, specify:

```
CHANNEL
```

```
                                CPU= 63 UIC=255 PDT=  46 DPR= 29   CHANNEL  T
11:43:02     CHANNEL     %CHP          CHANNEL     %CHP           CHANNEL     %CHP
            PATH TYPE    BUSY         PATH TYPE    BUSY          PATH TYPE    BUSY
             00   BL    15.98          01   BL     7.66           02   BL     9.63
             03   BL     0.73          04   BL     0.00           05   BL     4.23
             06   BL    17.61          07   BL    19.45           10   BL     4.23
             11   BL     7.56          12   BL     0.45           13   BL     7.45
             14   BL     0.18          15   BL    33.57           16   BL     0.00
             17   BL     3.02          20   BL     0.00           21   BL     3.25
             22   BL     0.00          23   BL     3.49           24   BL     0.00
             25   BL     0.00          26   BL     0.03           27   BL     0.00
             40   BL    17.77          41   BL     6.85           42   BL     0.00
             43   BL     0.68          44   BL     2.44           45   BL     0.00
             46   BL     0.00          47   BL    25.59           50   BY     0.00
             51   BL     7.09          52   BL     2.10           53   BL     1.86
             54   BL     0.10          55   BL    32.60           56   BL     0.00
             57   BL     5.91          60   BL    17.40           61   BL     6.61
             62   BL    12.47          63   BL     2.83           64   BL     0.00
             65   BL     0.00          66   BL    15.85           67   BL     0.00
```

Figure 4-17. Channel Path Activity (CHANNEL) Report

Field Heading	Meaning
CHANNEL PATH TYPE	The hexadecimal number that indicates the channel path identifier (CHPID) and either BY for byte multiplexor or BL for block multiplexor. For a 308x, 908x, or 3090 processor complex, RMF uses the I/O configuration data set (IOCDS) to find the channel path type. If RMF cannot read it, or if it has been updated so that the data might not apply to the present configuration, RMF does not report the channel path type; the field is blank. For a 4381 processor complex, RMF uses the I/O configuration data (IOCD) in the hardware system area (HSA) to retrieve the channel path type. If RMF cannot retrieve the configuration data after successive attempts, RMF does not report the channel path type; the field is blank.
% CHP BUSY	The percentage of time during the reporting interval when the channel path was busy. To calculate the value, RMF uses the sampling data the SRM collects in the channel path measurement table (CPMT). At the end of the reporting interval, the calculation is: % CH PATH BUSY = 100 x (SB/N) % CH PATH BUSY equals 100 times the number of SRM observations of channel path busy, divided by the number of samples taken. This field always contains a value that reflects the activity since the last report request.

Figure 4-18. Fields in Channel Path Activity (CHANNEL) Report

Domain Activity Report

The domain activity report includes the SRM data in the domain descriptor table at the time you request the report. The report allows you to obtain the same data as the system response to the DISPLAY operator command with the DMN option. The information in this report allows you to verify the effectiveness of the current domain definitions.

The heading of the report contains the name of the current installation performance specification (IPS) parmlib member, followed by the CPU, I/O, main storage, and SRB service coefficients, and the time when the report was requested.

Figure 4-19 shows a sample domain report. The field headings, their meanings, and the possible contents are described in Figure 4-20.

To request a domain activity report during a display session, specify:

```
DDMN
```

```
                               CPU= 15 UIC=255 PDT=   2 DPR= 14   DDMN     T
IPS = IEAIPSXI   CPU=10.0 IOC= 5.0 MSO= 3.0 SRB=10.0                 11:43:40
DMN  MIN MAX MPLT CMPL RUA INC NSW GOO OUTU  TWSR  TCTL    CIDX  RTOB
 0   999 999  999    4   5   0   4   0    0    1K   1F     0.00  0.0
 1     1  50    1    0   0   0   0   0    0  188    3D    98.14  0.0
 2     1  50    2    1   1   1   0   0    0  442    4D    91.25  0.0
 3    35  50   35    5   4   4   0   0    0    1K   1A   146.80  0.0
 4    10  50   10    0   1   0   0   0    0    1K   1A     0.00  0.0
 5    10  50   10    0   1   0   0   0    0    2K   1A     0.99  0.0
 6     2  50    2    0   0   0   0   0    0   31    2D    95.04  0.0
 7    17  50   21   21  20   3  18   0    0    1K   3D    82.52  0.0
 8     2  50    2    0   0   0   0   0    0    0    3D     0.00  0.0
 9     1   2    2    1   1   0   1   0    0  355    1F     1.00  0.0
10     0   0    0    0   0   0   0   0    0    0    4D     0.00  0.0
```

Figure 4-19. Domain Activity (DDMN) Report.

Field Heading	Meaning
DMN	The decimal identifier of the domain.
MIN	The minimum multiprogramming level specified for the domain in the IPS.
MAX	The maximum multiprogramming level specified for the domain in the IPS.
MPLT	The current multiprogramming level target the SRM has computed for the domain.
CMPL	The current actual multiprogramming level for the domain.
RUA	The average number of ready users, rounded off to the nearest integer.
INC	The current number of swappable in-storage users.
NSW	The current number of non-swappable in-storage users.
GOO	The current number of address spaces being swapped out.
OUTU	The current number of users on the out queue.
TWSR	The time-weighted average service rate for the domain. If the value exceeds 99999, the value is represented by nK, where n is the value divided by 1000. For example, 100,000 is displayed as 100K.
TCTL	The parameter value and method that the SRM is using for MPL target control. The parameter value appears in the first three positions; leading zeroes are suppressed. The method appears in the last position, indicated by one of the following: W indicates that the parameter value is a weighting factor as specified by the CNSTR keyword in an IPS or by the WT option in a SETDMN operator command. A indicates that the parameter value is a performance objective as specified by the AOBJ keyword in an IPS or a SETDMN operator command. The AOBJ value controls the average service rate to ready users in the domain. D indicates that the parameter value is a performance objective, as specified by the DOBJ keyword in an IPS or a SETDMN operator command. The DOBJ value controls the total service rate to the domain. F indicates that the parameter value is a performance objective as specified by the FWKL keyword in an IPS or SETDMN operator command. The FWKL value specifies a workload level that is used as a contention index for the domain.

CIDX	The contention index for the domain. The contention index is computed by the SRM in one of two ways:
	
	• If a domain weight is specified (indicated by a W in TCTL), the calculation is
	`CIDX = average ready users x weight / target MPL`
	
	• If a target control keyword is used (indicated by a A, D, or F in TCTL), the calculation is
	`CIDX = workload level`

RTOBT	The average elapsed time (response time) in tenths of a second for first period TSO transactions, if RTO was specified in the IPS for this IPL; otherwise, the field contains zero.

Figure 4-20. Fields in the Domain Activity (DDMN) Report

Enqueue Activity Report

The enqueue activity report enables you to track contention for or ownership of serially-reusable resources to which access is controlled by ENQ and DEQ macro instructions. The report is a "snapshot" report that describes the contention or ownership at the time RMF processes the request for the report.

Especially when invoked during a display session, the report can help you to determine on a realtime basis which resources and jobs are contributing to any bottlenecks caused by resource contention. The report describes only the contention caused by ENQ and DEQ macro instructions. For information about the contention caused by the RESERVE macro instruction, use the reserve activity report.

You can request the enqueue activity report as a summary report for all resources that had contention, as a detail report for all resources that had contention, as a detail report for a specific resource identified by name or as a report of resources held by a specific system, whether or not there is contention. The summary report includes all resources for which there is contention and describes the number of tasks that own each resource and the number of tasks waiting for the resource. The detail report includes all resources for which there is contention and identifies by jobname, system, and address space identifier the jobs that own the resource and the jobs that are waiting for the resource. When you request the report for a specific resource by major name or major and minor name, the report includes detail data for the resource or group of resources requested.

The report on the resources held by a specific system in a global resource serialization complex identifies either all the resources held or just the exclusively-held resources. This report is useful when attempting to recover an inactive system in a global resource serialization complex. You can request this report from an active system in the complex and determine from the report the resources that the inactive system held.

Figure 4-21 shows a sample summary report; Figure 4-22 shows a sample resource report for a specific system. Figure 4-23 shows a sample detail report. The field headings, meanings, and contents for all reports are described in Figure 4-24.

To request an enqueue activity report during a display session, specify:

```
SENQ [{S}][,majorname][,minorname]
      {D}
```



```
                                     CPU= 78 UIC=255 PDT=  44 DPR= 92   SENQ     T
11:44:26 TSK  TSK  TSK MAJOR NAME
         OWN  WTE  WTS   MINOR NAME
                       SYSDSN
           1    1    0   D75TYT1.IATCNNJ.P10.ASM(SYSS)
           1    0    1   GALER.PA21100.PTMPRINT(SYSS)
```

Figure 4-21. Enqueue Contention Activity (SENQ) Summary Report

```
                                     CPU= 78 UIC=255 PDT=  44 DPR= 92   SENQ     T
11:44:26                           MAJOR NAME
JOBNAME  SYSTEM  ASID  REQ   MINOR NAME
                           ADMPRNTQ
ADMPRINT M303     185   EO   BACKGROUND DRIVER EXECUTING.(SYSS)
                           ARCGPA
HSM      M303      13   SO   ARCMCDS  (SYSS)
                           BLXDASDS
D10MCW1  M303     116   SO   D83INFO.BLGPNLS                          (SYSS)
                           CLR12.UN
D32VJF1  M303      88   SO   IUSER                                    (SYSS)
                           GIMSMP
D31BAR1  M303     217   EO   C87JPLX.SMPCSI.CSI                       (SYSS)
                           SPFDSN
D31BAR1  M303      35   EO   D31BAR1.IAREND.P06.PLS                   (SYSS)
D75DVB1  M303      41   EO   D75DVB1.IATOSDR.W1A.ASM                  (SYSS)
D96JRF1  M303     163   EO   D96JRF1.IEFDB413.PS1.PLS                 (SYSS)
```

Figure 4-22. Enqueue Activity (SENQ) Global Resource Serialization Complex Resource Report

```
                                    CPU= 55 UIC=255 PDT=  42 DPR= 13   SENQ     T
11:48:44                     MAJOR NAME
JOBNAME  SYSTEM  ASID REQ     MINOR NAME
                             SYSDSN
D75TYT1  AQXI     136 EO   D75TYT1.IATCNNJ.P10.ASM(SYSS)
D75TYT1A AQXA      29 EW
GALER    AQXI      62 EO   GALER.PA21100.PTMPRINT(SYSS)
GALERH   AQXA      30 SW
```

Figure 4-23. Enqueue Contention Activity (SENQ) Detail Report

Field Heading	Meaning
TSK OWN	(Summary report only.) The number of tasks that currently own the resource.
TSK WTE	(Summary report only.) The number of tasks that currently are waiting for exclusive use of the resource.
TSK WTS	(Summary report only.) The number of tasks that currently are waiting for shared use of the resource.
JOBNAME	(Detail and specific system reports only.) The name of the job that has requested use of the resource.
SYSTEM	(Detail and specific system reports only.) The identifier of the system on which the job that owns or requests the resource is running.
ASID	(Detail and specific system reports only.) The address space identifier of the job that has requested use of the resource.
REQ	(Detail and specific system reports only.) A two-character field that describes the request. The first character indicates the type of the request and can be either: E - indicating that the request was for exclusive use of the resource S - indicating that the request was for shared use of the resource The second character indicates the status of the request and can be either: O - indicating that the requestor owns the resource W - indicating that the requestor is waiting for the resource
MAJORNAME MINORNAME	The name and scope of the resource. The major name, which corresponds to the *qname* field in the ENQ and DEQ macro instructions, is one to eight characters in length; it is aligned under the MAJORNAME heading. The minor name, which corresponds to the qname field in the ENQ and DEQ macro instructions, can be from 1 to 255 characters in length. However, only 44 characters can appear in the report. When a minor name exceeds 44 characters, it is truncated in the report, and an asterisk (*) following the scope indicates that the name has been truncated. If the minor name contains unprintable characters, RMF reports in in the form 'name', where *name* appears as up to 44 hexadecimal digits. Each minor name is aligned under the MINORNAME heading. RMF recognizes only 44 characters. Therefore if two minor names (both with the same major name) are longer than 44 characters and differ only beyond the forty-fourth character, RMF cannot distinguish between them. The scope of the resource follows the minor name. A resource with a scope of 'SYSTEMS' is followed by (SYSS). A resource with a scope of 'SYSTEM' is followed by (STEP).

Figure 4-24. Fields in the Enqueue Contention Activity (SENQ) Report

I/O Device Activity Report

The I/O device activity report gives information on I/O device use for all online devices you request either by device class, by device number, or by volume serial number. You can request the I/O device activity report either as a row report, which consists of one data line for one device, or as a table report, which consists of a variable number of data lines for one or more devices. Either form of the report can help you identify and correct bottlenecks caused by a particular device.

The report, like the Monitor I session report, can help you to analyze device performance, identify bottlenecks, and resolve obstacles that prevent efficient use of the resource. Requesting the report during a display session enables you, for example, to track the device use on a realtime basis. You can get a timely picture of device use or track a specific critical device on a realtime basis, thus making it possible to take any required immediate corrective action.

The information in the report can, depending on your choice, show either the cumulative total since the beginning of the Monitor I interval or the "delta" — the change since the previous request for the report. To request delta mode — the change in device activity since the last request for the report — use the D display command.

The report is based on both hardware measurements and data collected during a Monitor I session; therefore, a Monitor I session must be active when you issue your request. If no data is available, RMF issues a descriptive message.

Because the data comes from both hardware measurements and Monitor I session measurements, the data required to report some or all of the fields might be invalid or unavailable. A field based on data that is unavailable or invalid contains dashes (---). The fields that RMF might not be able to report and the possible reasons for the unavailable or invalid data are:

field name	not reported when
ACTIV RATE RESP TIME IOSQ TIME DB DELAY PEND TIME DISC TIME CONN TIME %DEV UTIL	Hardware measurements are not available either because there are not enough CMBs or because the channel measurement facility is disabled.
RESP TIME IOSQ TIME CUB DELAY PEND TIME DISC TIME CON TIME %DEV UTIL	The device is attached to a byte multiplexor channel. Byte multiplexor channels collect only activity rate data; that is, the only hardware measurement available is the start subchannel count (SSCH).
LCU	RMF was unable to read the IOCDS (308x, 908x, or 3090 processor) or the IOCD in the HSA (4381 processor).
ALL	The device moved online or offline during the Monitor I interval and total mode was requested.

STG GRP	Reported as **CHGD** when a device is added or deleted during the report interval.

If an '*' appears immediately to the right of a field, a hardware measurement timer overflow has occurred. See "Timer Overflow" in Chapter 3 for an explanation of this condition.

To evaluate the data, you need to understand what a reporting period is and how it relates to the Monitor I interval.

The I= field at the right of each report shows the percentage of the Monitor I interval that has elapsed when RMF generates the Monitor II session report. The field is calculated as follows:

$$I = \frac{\text{number of samples taken x cycle time}}{\text{Monitor I interval length}} \times 100$$

For a report that reflects the total device activity (delta mode is off), the reporting period is the time that has elapsed from the start of the Monitor I interval to the time when you request the report. The maximum reporting period is one Monitor I interval. When a Monitor II report covers a complete Monitor I interval, the I= field contains two asterisks (I=**).

For a report that reflects the change in device activity (delta mode is on), the reporting interval is the period between any two consecutive Monitor II requests. Your initial request establishes the base for reporting the activity but does not generate a report. After processing your initial request, RMF prompts you to press the enter key. Pressing the enter key causes RMF to generate the report. For your initial report, the reporting interval is the period between your initial request and the time you press the enter key. If you continue to request the report, the reporting interval is the time between each report request, and the data reflects the change in activity since your last request. If a Monitor I interval expires between two consecutive Monitor II requests, no data is reported and RMF prompts you to press the enter key.

The I= field at the right of the report equals the percentage of the interval that is represented by the data; thus, for your initial request, the I= field equals the percentage of the interval that expires between your initial request and the time you press the enter key. For all subsequent requests, the I= field equals the percentage of the interval that expires.

You can request the I/O device activity report either as a row report, which consists of one data line for one device, or as a table report, which consists of a variable number of data lines for one or more devices. The DEV report is sorted by LCU, unless you specify the storage group (SG) option. The SG option causes the DEV report to be sorted by device numbers within storage groups. The storage groups appear on the DEVV report only when the device has been assigned to a storage group. SG is not a valid option for DEVV. Figure 4-25 shows a sample I/O device row report. To request an I/O device activity row report during a display session, specify:

```
DEVV  { VOLSER (volid)          }
      { NUMBER (device-number)  }
```

Figure 4-26 shows a sample I/O device activity table report. To request an I/O device activity table report during a display session, specify:

```
DEV  {device-class                }
     {VOLSER (volid,...)          }
     {NUMBER (device-number)      }
     {SG (storage group name)     }
```

When you request the report during a display session, the data line for any device that is more than 30% utilized is highlighted.

The field headings, their meanings, and possible contents, which are the same for both the row report and the table report, are described in Figure 4-27.

```
                        MIG=*** CPU=102 UIC=174 PDT=  23 DPR=  1   DEVV      T
SG              DEV     ACTIV RESP IOSQ --DELAY-- PEND DISC CONN  %DEV %DEV
  TIME    VOLSER NUM LCU RATE  TIME TIME  CUB  DB  TIME TIME TIME  UTIL RESV I%
14:48:29 RAC001 631 027 24.031    9    3 0.0  3.0  3.3  1.4  1.8 17.39 17.2 61
14:48:30 RAC001 631 027 24.027    9    3 0.0  3.0  3.3  1.4  1.8 17.38 17.2 61
14:48:35 RAC001 631 027 24.048    9    3 0.0  3.0  3.3  1.4  1.8 17.44 17.2 62
14:48:36 RAC001 631 027 24.059    9    3 0.0  3.0  3.3  1.4  1.8 17.45 17.2 62
14:48:37 RAC001 631 027 24.065    9    3 0.0  3.0  3.3  1.4  1.8 17.44 17.2 62
14:48:38 RAC001 631 027 24.068    9    3 0.0  3.0  3.3  1.4  1.8 17.44 17.2 62
14:48:39 RAC001 631 027 24.070    9    3 0.0  3.0  3.3  1.4  1.8 17.43 17.2 62
14:48:40 RAC001 631 027 24.074    9    3 0.0  3.0  3.3  1.4  1.8 17.43 17.1 62
14:48:41 RAC001 631 027 24.059    9    3 0.0  3.0  3.3  1.4  1.8 17.51 17.2 62
14:48:42 RAC001 631 027 24.058    9    3 0.0  3.0  3.3  1.4  1.8 17.60 17.3 62
14:48:43 RAC001 631 027 24.052    9    3 0.0  3.0  3.3  1.4  1.8 17.59 17.3 62
14:48:44 RAC001 631 027 24.036    9    3 0.0  3.0  3.3  1.4  1.8 17.57 17.3 62




--------------------
```

Figure 4-25. I/O Device Activity Row (DEVV) Report


```
                                  D I R E C T   A C C E S S   D E V I C E   A C T I V I T Y

                          MIG=*** CPU=104 UIC= 72 PDT= 140 DPR= 27   DEV      T
14:39:49        DEV       ACTIV RESP IOSQ --DELAY-- PEND DISC CONN  %DEV %DEV
STG GRP  VOLSER NUM LCU   RATE  TIME TIME  CUB  DB  TIME TIME TIME  UTIL RESV I%
         MIG001 2DE 00F  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         D24PK1 2DF 00F  0.711     7    0  0.0  1.0  1.3  0.9  5.1 0.420  0.0 32
         D83STC 320 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         D83STD 321 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         CS2212 322 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         D54323 323 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         D54324 324 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         D54325 325 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         CE313S 326 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         CE3131 327 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         CE3132 328 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         C84DMP 329 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         TSODXX 32A 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         TSODL1 32B 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         TXADL1 32C 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         TXADL2 32D 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         TXADL3 32E 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         TXADL4 32F 013  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         RF1805 340 014  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         D94RM1 341 014  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         RF1815 342 014  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         D41141 343 014  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         D94RF1 344 014  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         D94RF2 345 014  0.152    27    0  0.0  3.0  3.2 15.9  7.5 0.356  0.0 32
         D94RF3 346 014  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         RF1902 347 014  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
         Y43APC 348 014  0.000     0    0  0.0  0.0  0.0  0.0  0.0 0.000  0.0 32
```

Figure 4-26. I/O Device Activity Table (DEV) Report (308x, 4381, 3090)

Field Heading	Meaning
STG GRP	The one to eight character name that identifies the storage group to which the device belongs. For DEVV, this field is reported only when the volumes specified are members of a storage group.
VOLSER	The six-character volume serial number (for direct access and magnetic tape reports) of the volume mounted on the device at the end of the reporting interval.
DEV NUM	The unique three-digit hexadecimal device number that identifies a physical I/O device.
LCU	The three-digit hexadecimal number that identifies the logical control unit (LCU) to which the device belongs. The set of devices associated with an LCU measurement are not the same on all processors because the definition of an LCU is model-dependent. On a 308x or 908x processor, an LCU is the set of devices attached to the same physical control unit (or a group of physical control units with one or more devices in common.) Each physical control unit and each device can belong to only one LCU; they cannot be shared between LCUs. On a 4381 processor complex, an LCU is the set of devices attached to an identical set of physical control units. Each device belongs to only one LCU, however, a physical control unit may belong to more than one LCU. On a 3090 processor, an LCU is the set of devices attached to the same physical control unit (or group of control units that have one or more devices in common). As with the 308x processor, each device belongs to only one LCU, but the I/O processor (IOP), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU. To find the LCU number, RMF must access the I/O configuration data set (IOCDS) in a 308x, 908x, or 3090 processor or the I/O configuration data (IOCD) in the hardware system area (HSA) in a 4381 processor. If RMF cannot read it, or if it has been updated so that the data might not apply to the present configuration, RMF does not report this field.
ACTIV RATE	The rate per second at which start subchannel (SSCH) instructions to the device completed successfully. In the line of data for the base exposure of a multiple exposure device, this field reflects the rate for the entire device. The data line for the base exposure is followed by data lines for each exposure. The value is calculated by dividing the total number of successful SSCH instructions by the number of seconds in the interval.
RESP TIME	The average number of milliseconds the device required to complete an I/O request. This value reflects the total hardware service time and the front end software queuing time involved for the average I/O request to the device. The channel measures active time, which starts at the acceptance of a SSCH instruction (indicated by a condition code (0) and ends at the acceptance of the channel end (primary status pending). It does not, however, include the time required to process the interruption. The IOS queue length is factored in to reflect the front end queuing time.
IOSQ TIME	The average number of milliseconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued. The calculation is: `AQT = (QC/S) / R` The average IOS queue time (AQT) equals the total count of requests found on the IOS queue at Monitor I samples (QC) divided by the number of Monitor I samples (s), divided by the device activity rate (R), described above.

Figure 4-27 (Part 1 of 3). Fields in the I/O Device Activity (DEV or DEVV)

Field Heading	Meaning
CUB DELAY	The average number of milliseconds of delay that I/O requests to this device encountered because the control unit was busy. If the device is shared at the control unit level, contention might be caused by the sharing system. If the device is not shared at the control unit level, the contention is the result of other activity to different devices over the same or alternate path serviced by the control unit. If the data is not valid or if the system is running on a 308x or 4381 processor, the field heading appears but a dash ('-') appears in the field in place of the data. The calculation is: `ACUB = CUBD / MEC` The average control unit busy delay (ACUB) equals the amount of control unit busy delay time (CUBD) divided by the measurement count count. In the line of data for the base exposure of a multiple exposure device, this field is the weighted average of the corresponding values for all exposures.
DB DELAY	The average number of milliseconds of delay that I/O requests to this device encountered because the device was busy. Device busy might mean that the volume is in use by another system, the device is reserved by another system, head of string busy condition caused the contention, or some combination of these conditions has occurred. If the data is not valid or if the system is running on a 308x or 4381 processor, the field heading appears but a dash ('-') appears in the field in place of the data. **In a PR/SM environment in LPAR mode,** this value is updated every 20 seconds. The calculation is: `ADB = DB / MEC` The average device busy delay (ADB) equals the amount of device busy delay time (DBD) divided by the measurement event count (MEC). In the line of data for the base exposure of a multiple exposure device, this field is the weighted average of the corresponding values for all exposures.
PEND TIME	The average number of milliseconds an I/O request remains queued in the channel. This value reflects the time between acceptance of the SSCH function at the subchannel (SSCH-function pending) and acceptance of the first command associated with the SSCH function at the device (subchannel active). This value also includes the time waiting for an available channel path and control unit as well as the delay due to shared DASD contention. If the value is high, refer to the device LCU entry in the I/O queuing activity report for an indicator of the major cause of the delay. The calculation is: `PT = FPT / MEC` The average pending time (PT) equals the total function pending time for the device (FPT) during the measurement interval divided by the measurement event count (MEC).
DISC TIME	The average number of milliseconds the device was disconnected (not transferring data) while processing an SSCH instruction. Thus, this value reflects the time when the device was in use but not transferring data. It includes the overhead time when a device might disconnect to perform positioning functions such as SEEK/SET SECTOR, as well as any reconnection delay. The calculation is: `ADT = TDT / MEC` The average disconnect time (ADT) equals the total disconnect time for the device (TDT) during the interval divided by the measurement event count (MEC). RMF calculates the total disconnect time by adding the pending time and connect time for the device and subtracting the result from the active time.

Figure 4-27 (Part 2 of 3). Fields in the I/O Device Activity (DEV or DEVV)

Field Heading	Meaning
CONN TIME	The average number of milliseconds the device was connected to a channel path and actually transferring data between the device and main storage. Typically, this value, measures data transfer time but also includes the search time and the time needed to maintain channel path, control unit, and device connection. The calculation used is: `ACT = DCT / MEC` The average connect time (ACT) equals the total connect time for the device (DCT) during the measurement interval divided by the measurement event count (MEC).
% DEV UTIL	The percentage of time during the interval when the device was in use. This percentage includes both the time when the device was involved in I/O operations (connect and disconnect time) and the time when it was reserved but not involved in an I/O operation. The percentage reported represents the time during the interval when the device is "tied up" when it could not be used to service a request from another system. RMF uses hardware measurement values to compute the time when the device was involved in an I/O operation. It uses sampling techniques to estimate the time when the device was reserved but not actually involved in an I/O operation, based on the settings of two bits: UCBSTRT is set on when the device is processing an I/O request and off when there is no active I/O request. UCBRESVH is on when the device is reserved. When UCBSTRT is off and UCBRESVH is on, RMF counts the sample as indicating non-overlapped device reserved time; that is, the device is reserved but not actively performing an I/O operation. Some small portion of device busy (reserved) time is missed when the device is reserved and UCBSTRT is on but the I/O request is pending in the channel. The calculation used is: `DUP = 100 x (DCT+TDT/T) + (RS/S)` The device utilization percentage (DUP) equals 100 times the sum of the device (DCT+TDT) connect and disconnect time, divided by the length of the measurement interval in seconds (T) and the number of Monitor I samples when the device was reserved but not involved in an I/O operation (RS) divided by the total number of Monitor I samples (S). The % DEV UTIL field on a device activity report can exceed 100% for a device that is 100% utilized. This is because the device connect time from the channel measurement block is a longer time period than the RMF measurement interval, so it is possible that the value can be slightly higher than 100%.
% DEV RESV	The percentage of time during the interval when a shared device was reserved by the processor on which RMF was started. At each RMF cycle, RMF checks to see if a device is reserved, and a counter is kept of all such samples. At the end of the interval, the percentage is computed. The calculation used is: `P = 100 x DRS/S` The percentage of time during the interval when the device is reserved (P) equals 100 times the number of device-reserved samples (DRS) divided by the total number of samples (S).

Figure 4-27 (Part 3 of 3). Fields in the I/O Device Activity (DEV or DEVV)

I/O Queuing Activity Report

The I/O queuing activity report provides information, grouped by logical control unit, on the I/O configuration. The information includes activity rate, queue lengths, and percentages of time when one or more I/O components were busy, grouped by logical control unit (LCU). Information about the LCU is especially useful because the LCU is the focus of I/O configuration and path management measurements for a related group of I/O devices. The definition of a logical control unit (LCU) is model-dependent.

On a 308x or 908x processor, an LCU is the set of devices attached to the same physical control unit (or a group of physical control units with one or more devices in common.) Each device can belong to only one LCU.

On a 4381 processor complex, an LCU is the set of devices attached to an identical set of physical control units. A physical control unit can belong to more than one LCU, if the devices are configured asymmetrically. A device can only belong to one LCU.

On a 3090 processor, an LCU is the set of devices attached to the same physical control unit (or group of control units that have one or more devices in common). As with the 308x processor, each device belongs to only one LCU, but the I/O processor (IOP), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU of a 3090 processor.

This report provides information about the cause of performance problems associated with channel paths and devices. For example, if the device report shows an unusually long pending time, you should look, on the I/O queueing report, at the percentage of requests deferred for device busy and control unit busy for the logical control unit. The relationship between these values indicates the major reason for the long pending time. Requesting the report during a display session enables you to track I/O queuing on a realtime basis, which can help you to determine the best way to resolve a performance problem related to a logical control unit.

The report depends on data that the Monitor I session collects or obtains from the IOCDS (for a 308x/908x or 3090 processor complex) or the hardware system area (for a 4381 processor complex). The Monitor I gatherer gets a new set of model dependent data every second or every cycle, whichever time period is greater.

If hardware measurements are not available on the 308x or 4381 processors, or if the data is not reliable (indicated by a successive invalid sample count greater than zero), RMF does not provide model-dependent data generated by the hardware for the following fields:

- ACTIV RATE
- AVG Q LNGTH

  (NOTE: The contents of this field are always suppressed for the 4381 processor.)
- %REQ DEFER

If there is no channel path status data available, the field % ALL CH BUSY is left blank.

If, on a 3090 processor, the hardware measurements are not available, the channel measurement facility is not available, or there is a failure in the diagnose interface, RMF does not provide model-dependent data generated by the hardware for the following fields:

- CONTENTION RATE
- DELAY Q LNGTH
- % ALL CH BUSY
- CHPID TAKEN
- %CU BUSY

All invalid data items are marked by dashes (---) in the output display.

The I= field in the heading of each report shows the percentage of the Monitor I interval that has elapsed when RMF generates the Monitor II session report. The field is calculated as follows:

```
I= (number of samples taken x cycle time)  x  100
        Monitor I interval length
```

For a report that reflects the total device activity (DELTA mode off), the reporting period is the time that has elapsed from the start of the Monitor I interval to the time when you request the report. The maximum reporting period is one Monitor I interval. When a Monitor II report covers a complete Monitor I interval, the I= field contains two asterisks (I=**).

The values in the report can, depending on your choice, show either the cumulative total since the beginning of the Monitor I interval or the "delta" -- the change since the previous request for the report. To request delta mode — the change in activity since the last request for the report — use the D display command.

In delta mode your initial request establishes the base for reporting the activity but does not generate a report. After processing your initial request, RMF prompts you to press the ENTER key. Pressing the ENTER key causes RMF to generate the report. The reporting interval is the period between any two consecutive Monitor II requests. For your initial request, the reporting interval is the period of time between your initial request and the time you press the ENTER key. If you continue to request the report, the reporting interval is the time between each report request, and the data reflects the change in I/O queuing activity since your last request.

When a logical control unit has no activity during the interval, that logical control unit is omitted from the report for that interval. If channel paths were brought online or taken offline during the interval, data is still formatted, but only the channel paths and control units that were online and had some connection to a device or set of devices of the logical control unit (LCU) at the time the report was requested appear in the report.

To request an I/O queuing report during a display session, specify:

```
IOQUEUE (type)
```


Figure 4-28 shows a sample I/O queuing report for 308x, 908x, and 4381 processors. Note that the AVG Q LNGTH field heading appears in the report but that the field contents are suppressed for the 4381 processor. The field headings, their meanings, and the possible contents are described in Figure 4-29.

```
F                                   CPU= 29 UIC=255 PDT=    0 DPR= 16    IOQUEUE  T
11:55:08          I=83%                 ACTIV  AVG Q % ALL   % REQ  %REQ DEFER
CHANNEL PATHS     CONTROL UNITS  LCU  RATE   LNGTH CHBUSY  DEFER DEVBSY CUBUSY
 01,11,41,51     102,103         006  12.638  0.37   2.79  18.26   0.00  18.26
 01,11,41,51     104,105         008  12.164  0.32   2.79  18.12   0.00  18.12
 62,02           203,230         00B  21.075  0.00   1.73  11.01   7.74   3.27
 62,02           201,210         00C   4.165  0.00   1.73   7.78   2.12   5.66
 62,02           204             00D  19.033  0.01   1.73  11.91   1.85  10.06
 62,02           200             00E   0.565  0.00   1.73   1.85   1.15   0.69
 54,14           400,410         016   1.043  0.00   0.00  21.79  20.45   1.34
 54,14           402             017   0.013  0.00   0.00  25.93   0.00  25.93
 66,26           604,640         01B   0.084  0.00   0.00  22.42  22.42   0.00
 66,06           601,610         01C  71.624  0.00   2.66  11.67   8.36   3.30
 66,06           607             01F   0.722  0.00   2.66  12.40   4.35   8.05
 66,06           605             021   3.579  0.00   2.66   1.03   0.94   0.09
 66,06           603,630         022  14.727  0.00   2.66   5.77   3.36   2.41
 66,06           600,6FF         023  14.187  0.00   2.66  23.20  11.17  12.03
 47,07           700,710         025  18.403  0.01   5.31  16.50   7.35   9.15
 47              701             026  27.039  0.05  21.25  21.01   3.05  17.96
 47,07           703,730         028   7.581  0.00   5.31  12.53  11.19   1.33
 47,07           702,720         029   1.123  0.00   5.31   0.94   0.82   0.12
 47,07           704             02A  29.083  0.05   5.31  12.65   1.39  11.26
 53,21           801             02E   0.019  0.00   0.07   0.00   0.00   0.00
 61,57           805,813         030  28.670  0.00   0.46   1.07   0.99   0.08
```

Figure 4-28. I/O Queuing Activity (IOQUEUE) Report (308x, 908x, and 4381)

Field Heading	Meaning
CHANNEL PATHS	The two-digit hexadecimal channel path identifier (CHPID) of the online channel path attached to the physical control units in the logical control unit group. There can be up to four channel paths in a logical control unit. Only channel paths identified in the Monitor I report as ONLINE to the system and having connection to a device or group of devices of the LCU will be listed in the Monitor II report.
CONTROL UNITS	The three-digit hexadecimal identifier of each physical control unit associated with an online channel path in the logical control unit group. There can be up to four physical control units in a logical control unit.
LCU	The three-digit hexadecimal number that identifies the logical control unit (LCU). The set of devices associated with an LCU measurement are not the same on all processors because the definition of an LCU is model-dependent. On a 308x or 908x processor complex, an LCU is the logical representation of a physical control unit or a group of physical control units with one or more devices in common. Each physical control unit and each device can belong to only one LCU; they cannot be shared between LCUs. On a 4381 processor complex, an LCU is the logical representation of a set of devices attached to an identical set of physical control units. Each device belongs to only one LCU, however, a physical control unit may belong to more than one LCU. To find the LCU number, RMF must access the I/O configuration data set (IOCDS) in a 308x or 908x processor or the I/O configuration data (IOCD) in the hardware system area (HSA) in a 4381 processor. If RMF cannot read it, or if it has been updated so that the data might not apply to the present configuration, RMF terminates the I/O queuing activity report.
ACTIV RATE	The rate per second at which I/O requests were successfully selected from the logical control unit queue. When this rate is consistently relatively high, the physical control units in the logical control unit might be overloaded. The activity rate is calculated by dividing the total number of I/O requests queued during the interval by the number of seconds in the interval.
AVG Q LNGTH	This field appears only when I/O queuing activity is measured on a 308x/908x processor complex and is the average number of requests on the logical control unit queue. Each time a request is taken off the queue, RMF counts the number of requests remaining and adds it to an accumulator. At the end of the interval, RMF divides the accumulated number of requests by the number of times an I/O request was dequeued. The calculation used is: `AQL = ER / DR` The average queue length (AQL) for this interval equals the total number of enqueued requests (ER) divided by the number of dequeued requests (DR). For a 4381 processor, I/O requests are queued on a physical control unit basis within the channel subsystem. Because queue length information is not available from the channel subsystem, the AVG Q LNGTH field heading appears in the report, but the field contents are suppressed for a 4381 processor.

Figure 4-29 (Part 1 of 2). Fields in the I/O Queuing Activity (IOQUEUE) Report (308x, 908x, 4381)

Field Heading	Meaning
% ALL CH BUSY	The percentage of time during the measurement interval when all channel paths belonging to the logical control unit were busy at the same time. At each sampling cycle, RMF checks the results of the last Store Channel Path Status (STCPS) instruction that the SRM issued to determine whether all channel paths are busy. RMF performs this check only for those channel paths that are online at that sample. At the end of the interval, RMF computes the percentage. The calculation used is: `ACB = 100 x (SB/S)` The percentage of time during the interval when all channel paths in the logical control unit were busy (ACB) equals 100 times the number of Monitor I samples that indicated all channel paths were busy (SB) divided by the number of Monitor I samples taken (S).
% REQ DEFER	The percentage of unsuccessful attempts to start I/O requests during the measurement interval. Requests are deferred mainly because the control unit is busy or the device is busy. Another reason for delay might be a disabled path. The calculation used is: `RDP = 100 x (RD/RA)` The percentage of requests deferred (RDP) equals 100 times the number of requests deferred (RD) divided by the number of initial attempts to select a request (RA). Note that RMF calculates the number of requests deferred (RD) by subtracting the successful attempts to select a request from the number of initial attempts to select a request. Note that this value shows no indication of the magnitude in time of the deferred request. Therefore, evaluate this value in light of the pending times of the devices in the logical control unit.

% REQ DEFER DEV BSY CU BUSY	The percentage of I/O requests deferred because the device was busy (DEV BUSY) or because the control unit was busy (CU BUSY). RMF uses a separate calculation for each percentage.
	To determine the percentage of requests deferred because the device was busy, the calculation is:
	`PDB = 100 x (RDB/R)`
	The percentage (PDB) equals 100 times the number of requests deferred because the device was busy (RDB) divided by the number of initial attempts to select a request (R).
	To determine the percentage of requests deferred because the control unit was busy, the calculation is:
	`PCB = 100 x (RCB/R)`
	The percentage (PCB) equals 100 times the number of requests deferred because the control unit was busy (RCB) divided by the number of initial attempts to select a request (R).

Figure 4-29 (Part 2 of 2). Fields in the I/O Queuing Activity (IOQUEUE) Report (308x, 908x, 4381)

Figure 4-30 shows a sample I/O queuing report for the 3090 processor. The report includes information about the effective use of channel paths for individual LCUs.

The field headings, their meanings, and the possible contents are described in Figure 4-31.

```
F                                   CPU= 29 UIC=255 PDT=   0 DPR= 16   IOQUEUE  TH
11:55:08          I=83%             CONTENTION   DEL Q    %ALL      CHPID   %CU
CHANNEL PATHS     CONTROL UNITS  LCU   RATE      LNGTH   CH BUSY   TAKEN    BUSY
 20              00A             008  0.000     0.00     0.00      0.00  33.33
 11              00B             009  0.000     0.00     0.00      0.00   0.00
 20              013             00A  0.000     0.00     0.00      0.00  50.00
 2A              014             00B  0.015     0.00     0.09      0.00   0.00
 01              009             00C  0.000     0.01     0.00      0.00  19.05
 0B              011             00D  0.000     0.00     0.00      0.00   9.09
 29              010             010  0.082     0.15     0.00      3.840 29.21
 2F                                                                4.302 18.35
 25              027                                               4.438 19.79
```

Figure 4-30. I/O Queuing Activity (IOQUEUE) Report 3090

Field Heading	Meaning
CHANNEL PATHS	The two-digit hexadecimal channel path identifier (CHPID) of the online channel path attached to the physical control units in the logical control unit group. There can be up to four channel paths in a logical control unit. Only channel paths identified in the Monitor I report as ONLINE to the system and having connection to a device or group of devices of the LCU appears in the Monitor II report.
CONTROL UNITS	The three-digit hexadecimal identifier of each physical control unit associated with an online channel path in the logical control unit group. There can be up to four physical control units in a logical control unit.
LCU	The three-digit hexadecimal number that identifies the logical control unit (LCU). The set of devices associated with an LCU measurement are not the same on all processors because the definition of an LCU is model-dependent. On a 3090 processor, an LCU is the set of devices attached to the same physical control unit or a group of physical control units with one or more devices in common. Each physical control unit and each device can belong to only one LCU; they cannot be shared between LCUs. To find the LCU number, RMF must access the I/O configuration data set (IOCDS) in a 3090 processor. If RMF cannot read it, or if it has been updated so that the data might not apply to the present configuration, RMF ends the I/O queuing activity report.
CONTENTION RATE	The rate per second at which the IOP places delayed I/O requests on the CU-HDR for this LCU. For the 3090 processor, the IOP places an I/O request on the CU-HDR when all paths to the subchannel are busy and at least one path to the control unit is busy. For devices with only one path, or for devices where multiple paths exist and the busy condition is immediately resolved, the IOP does not count the condition. The calculation used is: `CR = ER / SI` The contention rate (CR) equals the total number of enqueued requests (ER) divided by the number of seconds in the interval (SI).
DEL Q LNGTH	The average number of delayed requests on the control unit header (CU-HDR). Each time a request is enqueued from the CU-HDR, RMF counts the number of requests remaining on the queue and adds that number to the accumulator. At the end of the interval, RMF divides the total number of accumulated queued requests by the number of times a request was enqueued. The calculation used is: `DQL = (AL-ER) / ER` The delayed queue length (DQL) equals the accumulated queue length (AL) minus the total number of enqueued requests (ER) divided by the total number of enqueued requests (ER).

Figure 4-31 (Part 1 of 2). Fields in the I/O Queuing Activity (IOQUEUE) Report (3090)

Field Heading	Meaning
% ALL CH BUSY	The percentage of time during the measurement interval when all channel paths belonging to the logical control unit were busy at the same time. Unlike the corresponding field in the report for the 308x and 4381 processors, this field is not based on sampling. The data is a mathematical probability derived from multiplying the percent busy of each channel path involved. Only channel paths that are both online to the system and connected to a device are included in the the calculation. If a channel path was online to the system but was not connected to any device or group of devices of the LCU, it is omitted from the calculation. The calculation used is: `ACB = CHPID0 x CHPID1 x CHPID2 x CHPID3` The percentage of time during the interval when all channel paths in the logical control unit were busy (ACB) equals the percent busy of each channel path involved (CHPID0, CHPID1, CHPID2, CHPID3).
CHPID TAKEN	The rate at which I/O requests to devices of this LCU are satisfied by each CHPID during the interval. By reviewing the rate at which each channel path of the LCU satisfies I/O requests, you can see how evenly the work requests are distributed among the available paths and how effectively those paths are arranged for the LCU. The calculation used is: `CPT = TO / SI` The rate at which I/O requests to devices of the LCU are satisfied by each CHPID (CPT) equals the total I/O operations accepted on that path (TO) divided by the seconds in the interval (SI).
% CU BUSY	Each CHPID of the LCU is a measure of how the control unit contention is distributed. There is a subtle difference in interpretation of this field when viewed against the existing field of %REQ DEFER - CUB which applies to the 308X and 4381 environments. The %CU BUSY field in 3090 shows the relationship for each channel path of the LCU, between requests deferred due to control unit busy and total successful requests serviced by that path. In 308X/4381, the %REF DEFER - CUB is a single value for the LCU which describes the relationship between the number of control unit busy conditions encountered and the total attempts (successful and unsuccessful), to service requests for devices of this LCU, via any channel path. The calculation used is: `PCB = 100 * (CUB / (CUB + CPT))` The percentage (PCB) equals 100 times the number of requests deferred because the control unit was busy (CUB) divided by the total I/O operations accepted on that path of the LCU (CPT) plus the number of requests deferred because the control unit was busy.

Figure 4-31 (Part 2 of 2). Fields in the I/O Queuing Activity (IOQUEUE) Report (3090)

Page/Swap Data Set Activity Report

The page/swap data set activity report provides information on either page data set activity or swap data set activity. A Monitor I session measuring page/swap data set activity must be active when you request the report.

Your initial request establishes the base for reporting the activity but does not generate a report. After processing your initial request, RMF prompts you to press the enter key. Pressing the enter key causes RMF to generate the report. The reporting interval is the period between any two consecutive Monitor II requests. For your initial request, the reporting interval is the period between your initial request and the time you press the enter key. If you continue to request the report, the reporting interval is the time between each report request, and the data reflects the change in page/swap data set activity since your last request.

Dashes (---) in the data fields indicate that RMF could not provide a value because the page or swap device has been varied online during the reporting interval.

The page/swap data set activity report can help you to determine whether the optimum size has been allocated for each page or swap data set. When you request the report during a display session, the data line for any fixed head paging device is highlighted.

Figure 4-32 shows a sample page data set activity report. Figure 4-33 shows a sample swap data set activity report. The field headings, their meanings, and the possible contents are described in Figure 4-34.

To request a swap data set activity report during a display session, specify:

```
PGSP SWAP
```

To request a page data set activity report during a display session, specify:

```
PGSP PAGE
```

```
                                   CPU= 57 UIC=255 PDT=  30 DPR= 23   PGSP      T
S VOLUME DEV  DEV   %SLOTS   PAGE   I/O REQ AVG PAGES        11:56:06
T SERIAL NUM TYPE   IN USE TRAN TIME  RATE   PER I/O V  DATA SET NAME
P RES84B 861 3380   41.89  ------   -------    0.500    SYS1.PLPAJ
C RES84B 861 3380    1.17  ------   -------    0.000    SYS1.COMMONJ
L PAG1C1 1C1 3351P  39.94  ------   -------    1.923 N  SYS1.PAG1C1
L PAG1C3 1C3 3351P  38.59  ------   -------    1.853 N  SYS1.PAG1C3
L PAG181 181 3351P  37.36  ------   -------    1.797 N  SYS1.PAG181
L PAG183 183 3351P  37.11  ------   -------    1.834 N  SYS1.PAG183
L PAGE90 6B5 3380    6.47  ------   -------    3.561 Y  SYS1.PAGE90V
L PAG862 862 3380    8.71  ------   -------    3.569 Y  SYS1.PAG862
```

Figure 4-32. Page Data Set Activity (PGSP) Report


```
                                   CPU= 64 UIC=255 PDT=  18 DPR= 23   PGSP      T
S VOLUME DEV  DEV    % SETS   AVG     I/O REQ                 11:56:56
T SERIAL NUM TYPE    IN USE SERV TIME  RATE              DATA SET NAME
S PAG1C2 1C2 3351P  12.93  ------   -------             SYS1.SWP1C2
S PAG180 180 3351P  12.47  ------   -------             SYS1.SWP180
S PAG182 182 3351P  12.40  ------   -------             SYS1.SWP182
S PAG1C0 1C0 3351P  12.87  ------   -------             SYS1.SWP1C0
```

Figure 4-33. Swap Data Set Activity (PGSP) Report

Field Heading	Meaning
S T	This field contains a one-letter identifier that identifies the type of paging space. The identifiers are: P - PLPA C - common D - duplex L - local S - swap
VOLUME SERIAL	This field indicates the volume serial number of the volume on which the paging data set resides.
DEV NUM	This field indicates the number of the devices on which the paging data set resides.
DEV TYPE	This field indicates the device type on which the paging data set resides.
% SLOTS IN USE	This field indicates the percentage of the slots in the paging data set that are in use. This field appears only in the page data set activity report. When you request the report, RMF calculates the percentage from the Monitor I sampling values. The calculation used is: $P = \left(\frac{N-A}{N}\right) x\ 100$ The percentage of slots in use equals the number of slots in the data set (N) minus the number of available slots (A) divided by the number of slots in the data set.
% SETS IN USE	This field indicates the percentage of sets in the swap data sets that are in use. This field appears only in the swap data set activity report. When you request the report, RMF calculates the percentage from the current Monitor I session sampling values. The calculation is: $P = \left(\frac{N-A}{N}\right) x\ 100$ The percentage of sets in use equals the number of sets in each swap data set (N) minus the number of available sets (A) divided by the number of sets in the data set.
AVG SRV TIME	This field appears only in the swap data set report. It indicates the average number of seconds required to complete each I/O request for a data set. When you request the report, RMF calculates the average service time. The calculation used is: T = (N/S) x (L/R) The average service time (T) equals the number of samples when the data set is in use (N) divided by the total number of samples (S) multiplied by the current length of the reporting interval (L) divided by the total number of Start I/O instructions issued to the data set during the interval (R). This field always contains a value that reflects the activity since the last report request.

Figure 4-34 (Part 1 of 2). Fields in the Page Swap Activity (PGSP) Report

Field Heading	Meaning
PAGE TRAN TIME	The field appears only in the page data set report. It indicates the page transfer time in seconds. When you request the report, RMF calculates the value from the current Monitor I sampling values. The calculation used is: $T = \left(\frac{N \times I}{S} \right) / R$ The average page transfer time (T) equals the number of samples when the data set was in use (N) multiplied by the length of the reporting interval (I) divided by the total number of samples (S), with the result being divided by the total number of pages transferred (R). This field always contains a value that reflects the activity since the last report request.
I/O REQ RATE	This field indicates the number of I/O requests per second for the data set made between the beginning of the interval and the time you request the report. This field always contains a value that reflects the activity since the last report request.
AVG PAGES PER I/O	This field appears only in the page data set report. It indicates the average number of pages that were transferred to or from the page data set. This field always contains a value that reflects the activity since the last report request.
V	This field indicates whether or not the local paging data set accepts VIO pages. The symbols are: Y - indicates VIO pages are accepted N - indicates VIO pages are not accepted
DATA SET NAME	This field indicates the name of the page or swap data set being monitored. If a page data set name or a swap data set name is longer than 23 characters, it is truncated in the report to 22 characters *followed* by an asterisk (*). If a data set has bad slots, the data set name is *preceded* by an asterisk (*). When the operating system detects errors in a data set that prevent its further use, RMF can no longer monitor the data set. RMF indicates that monitoring is terminated by *preceding* the data set name with two asterisks (**).

Figure 4-34 (Part 2 of 2). Fields in the Page Swap Activity (PGSP) Report

Paging Report

The Monitor II session paging activity report presents overview information on system paging activity. Each report consists of one line of data that gives a "snapshot" view of system paging activity at the time the report was requested. When you invoke the report repetitively, you can build a table showing the differences over a period of time.

All of the fields in the report except the time (TIME), length of the available frame queue (AFC), highest UIC (HI UIC), and the average of the highest UIC (A HI UIC) reflect rates that show the activity since the last request for the report to enable you to see more clearly the differences with each time interval.

Figure 4-35 shows how you can repeat requests for the report to build a table of system paging activity. The field headings, their meanings, and the possible contents are described in Figure 4-36.

To request a paging report during a display session specify:

```
SPAG
```

```
                                        CPU= 50 UIC=255 PDT=  25 DPR= 14    SPAG    T
          LPA CSA CSA SWP  PGS-SWPD PRIVATE VIO TAR         HI   ES MIG  ESF  MIG
  TIME    IN  IN  OUT OUT    IN  OUT IN  OUT I+O CWS   AFC UIC  RTE AGE  AVL  RTE
09:01:19 --- --- --- --- ---- ---- --- --- ---   0  4130 255
09:01:22 0.0 0.0 0.0 0.3 48.0 24.0 1.0 1.7 0.0   0  3982 255
09:01:26 0.0 0.0 0.0 0.8 59.8 60.5 0.3  16 1.0   0  4021 255
09:01:31 0.0 0.0 0.0 0.6 43.2 57.4 0.2 7.4 0.0   0  4022 255
09:01:33 0.0 0.0 0.0 1.0 36.0 39.5  12 2.5 0.0   0  4242 255
09:01:38 0.0 0.0 0.0 0.4 60.8 46.4  92  28 0.0   0  4124 255
09:01:42 0.0 0.0 0.0 0.3 37.5  0.0 1.3 1.0 0.0   0  3953 255
09:01:48 0.0 0.0 0.0 0.5 51.8 63.7 0.7  14 1.2   0  3802 255
09:01:52 0.0 0.0 0.0 1.0 88.3 46.3 5.0  20 0.8   0  3595 255
09:01:53 0.0 0.0 0.0 1.0  0.0  111 7.0 4.0 0.0   0  3770 255
09:01:54 0.0 0.0 0.0 0.0  0.0  0.0  11 0.0 0.0   0  3669 255
09:01:55 0.0 0.0 0.0 2.0  0.0  144  12  74 0.0   0  3654 255
09:01:58 0.0 0.0 0.0 0.7 26.3 50.3 1.7 7.0 0.3   0  3003 255
09:01:59 0.0 0.0 0.0 1.0 72.0  167  23  16 0.0   0  3076 255
09:02:00 0.0 0.0 0.0 1.0  0.0  0.0  10  17 5.0   0  3456 255
```

Figure 4-35. Monitor II Paging (SPAG) Report

Field Heading	Meaning
TIME	Time the report was requested, in the form hh.mm.ss.
LPA IN	The rate (number per second) of LPA pages paged in; after the initial request, this field represents the rate since the previous report.
CSA IN	The rate (number per second) of CSA pages paged in; after the initial requests, this field represents the rate since the previous report.
CSA OUT	The rate (number per second) of CSA pages paged out; after the initial request, this field represents the rate since the previous report.
SWP OUT	The rate (number per second) of successful swap-outs.
PGS-SWPD IN	The rate (number per second) of pages swapped in: after the initial request, this field represents the rate since the previous report.
PGS-SWPD OUT	The rate (number per second) of pages swapped out; after the initial request, this field represents the rate since the previous report.
PRIVATE IN	The rate (number per second) of private area (VIO + non-VIO) pages paged in, after the initial request, this field represents the rate since the previous report.
PRIVATE OUT	The rate (number per second) of private area (VIO + non-VIO) pages paged out, after the initial request, this field represents the rate since the previous report.
VIO I+O	The rate (number per second) of VIO pages paged in and paged out; after the initial request, this field represents the rate since the previous report.
TAR CWS	One to three digit number indicating the target working set size for the common area.
AFC	The total number of frames currently on all available frame queues.
HI UIC	The highest unreferenced interval count (UIC).
ES RTE	The rate (number per second) of pages sent to extended storage; after the initial request, this field represents the rate since the previous report. In a system without extended storage, the field heading appears in the report, but the field is blank.
MIG AGE	The length of time a page resides on extended storage before it migrates to auxiliary storage; this field provides a snapshot of the migration age taken at the last sample. In a system without extended storage, the field heading appears in the report, but the field is blank.
ESF AVL	The number of extended storage frames currently available and not in use. In a system without extended storage, the field heading appears in the report, but the field is blank.
MIG RTE	The rate (number per second) of page migration from extended storage to auxiliary storage; after the initial request, this field represents the rate since the previous report. In a system without extended storage, the field heading appears in the report, but the field is blank.

Figure 4-36. Fields in the Monitor II Paging (SPAG) Report

Real Storage/Processor/SRM Report

The real storage/processor/SRM report presents a one-line summary of the current utilization of real storage, the processor, and SRM facilities. It is a "snapshot" of system activity at the time that the report is requested. When you repeat the requests for the report, you can build a table showing the differences over a period of time. Figure 4-37 shows how repeated requests for the report can build such a table.

This report can be extremely helpful in alerting you to any abnormal situations, thus enabling you to obtain detailed reports on a timely basis that can give the additional information you need to fully analyze the problem.

The field headings, their meanings, and the possible contents are described in Figure 4-38.

To request a real storage/processor/SRM report during a display session, specify:

SRCS

```
                                  CPU= 49 UIC=255 PDT=  25 DPR= 47    SRCS    T
                HI SQA  LPA LPA  CSA L+C  PRV LSQA CPU  IN OUT OUT OUT ASM
  TIME     AFC UIC   F    F  FF    F  FF   FF   FF UTL   Q LOG  RQ  WQ   Q
09:02:30 3549 255 618 1097  49  748 205  346 3757  22  28 135   0  31  14
09:02:33 3516 255 618 1097  49  749 205  339 3764  13  25 138   0  31   8
09:02:35 3418 255 618 1097  49  748 205  344 3810  30  27 138   0  29   8
09:02:39 3566 255 619 1098  49  748 205  352 3783  31  26 138   0  31  10
09:02:41 3054 255 619 1098  49  749 205  338 3807  53  28 137   0  30  10
09:02:49 3216 255 619 1098  49  749 205  353 3811  87  31 134   0  31   3
09:02:59 3785 255 621 1098  49  749 205  349 3832  31  28 137   0  30  54
09:03:08 4559 255 622 1098  49  749 205  340 3806  29  26 138   0  31  17
09:03:15 3865 255 622 1098  49  751 205  340 3853  80  33 133   1  30  14
09:03:20 4119 255 622 1098  49  750 205  346 3875  32  29 139   0  29   6
09:03:29 3981 255 622 1098  49  750 205  338 3811  65  30 134   1  31   5
09:03:35 3691 255 622 1098  49  750 205  345 3835  49  27 139   0  30   7
```

Figure 4-37. Real Storage/Processor/SRM (SRCS) Report

Field Heading	Meaning
TIME	The time the report was requested, in the form hh.mm.ss.
AFC	The average length of the available frame queue.
HI UIC	The highest unreferenced interval count (UIC).
SQA F	The total number of SQA frames.
LPA F	The total number of LPA frames.
LPA FF	The total number of fixed LPA frames.
CSA F	The total number of CSA frames.
L+C FF	The total number of fixed LPA and CSA frames.
PRV FF	The total number of private non-LSQA fixed frames.
LSQA FF	The total number of private LSQA fixed frames.
CPU UTL	A one to three digit number indicating the current processor utilization. T value for all processors currently online. The calculation is: $\text{CPU UTL} = \dfrac{\text{Interval Time} - \dfrac{\text{Sum of Wait Times}}{\text{\# of Online Processors}}}{\text{Interval Time}} \times 100$ Values up to and including 128% can appear in this field. Values greater than 100% indicate the number of users ready and waiting to be dispatched. (110% indicates that 10 users are waiting to be dispatched) Although 128% is the highest number that can appear, more than 28 users may be waiting to be dispatched. **In a PR/SM environment in LPAR mode,** the following calculations are used to determine the current average processor utilization value: • For a dedicated partition: The above calculation and limitations apply. • For a non-dedicated partition: This value depends on Monitor I measurements. If Monitor I is not measuring CPU activity, a value is unavailable and "***" will appear in this field. Values up to 100% can appear in this field and are updated every five seconds. with Wait Completion YES: $\text{CPU UTL} = \dfrac{\text{Total Processor Dispatch Time} - \text{Sum of Wait Times}}{\text{Interval Time} \times \text{\# of online processors assigned to partition}} \times 100$ with Wait Completion NO: $\text{CPU UTL} = \dfrac{\text{Total Processor Dispatch Time}}{\text{Interval Time} \times \text{\# of online processors assigned to partition}} \times 100$ Processor Dispatch Time (PDT), Wait Time (WT), and Wait Completion (WC) are defined under the partition data report field descriptions in chapter 3.

Field Heading	Meaning
IN Q	The current length of the SRM in queue.
OUT LOG	The current number of address spaces that are logically swapped out.
OUT RQ	The current length of the SRM out ready queue.
OUT WQ	The current length of the SRM out wait queue.
ASM Q	The current SRM view of the length of the ASM queue.

Figure 4-38 (Part 2 of 2). Fields in the Real Storage/Processor/SRM (SRCS) Report

Reserve Activity Report

The reserve activity report enables you to track RESERVE macro instructions issued to reserve a shared direct access device (shared DASD) for use by a particular system. The report is a 'snapshot' report that describes the status of all RESERVE requests outstanding at the time RMF processes the request for the report.

Figure 4-39 shows a sample reserve activity report requested for all devices. The field headings, meanings, and contents are described in Figure 4-40.

To request a reserve activity report during a display session, specify:

```
SENQR [ ALLVSER ]
      [ volid   ]
```

```
                                CPU= 49 UIC=255 PDT=  13 DPR= 13   SENQR    T
11:58:17        SYSTEM ENQUEUE RESERVE REPORT
JOBNAME ASID SYSTEM  REQ VOLUME DEV RSV MAJOR    MINOR
#2030429  25 AQXI     SO POK811 8E3 CNV CLRLOG00 CLRVSAM.CONSOL.BULK0001        *
#2030429  25 AQXI     EO RMFPAC 6C1  ON SYSIGGV2 X'C3D3D9E5E2C1D44BC3D6D5E2D6'*
JES2      22 AQXI     EO CHKPT1 713 OFF SYSZJES2 CHKPT1SYS1.HASPCKPT            *
LAURILY  236 AQXI     SO RAC001 631 OFF SYSZRACF SYS1.RACF
D24AJD1J  40 AQXI     EO POK805 A15 CNV UNIENQ00 CLRVSAM.JOBLOG.UNIBCKUP        *
G50KAP1A  32 AQXA     EW
#2030428  32 AQXI     SW POK805 A15 CNV
GINNIA   182 AQXI     SO POK805 A15 CNV          POKLCG.COPY.SRVPLAN            *
F22TEMPP  26 AQXA     SO                         POKLCG.COPY.UNIPLAN            *
```

Figure 4-39. Reserve Activity (SENQR) Report

Field Heading	Meaning
JOBNAME	The name of the job that issued the RESERVE macro instruction for the device identified under DEV.
SYSTEM	The identifier of the system on which the job that owns or requests the resource is running.
ASID	The address space identifier of the job that issued the RESERVE macro instruction for the device identified under DEV.
REQ	A two-character field that describes the request. The first character indicates the type of the request and can be either: E - indicating that the request was for exclusive use of the device S - indicating that the request was for shared use of the device The second character indicates the status of the request and can be either: O - indicating that the requestor owns the device W - indicating that the requestor is waiting for the device
VOLUME	The one to six character volume serial of the volume mounted on the device identified under DEV. If reserves are issued on systems other than the one on which you request the report, the field is blank.
DEV	The three-character address of the device for which the RESERVE macro instruction was issued. If reserves are issued on systems other than the one on which you request the report, the field is blank.
RSV	An indicator of the reserve status of the device. When the field contains ON, the device is reserved by the processor on which RMF is running. When the field contains OFF, the device is reserved by another processor. If reserves are issued on systems other than the one on which you request the report, the field is blank.
MAJOR MINOR	The name used to control access to the device by means of RESERVE macro instructions. The major name, which corresponds to the *qname* field in the RESERVE macro instruction, is one to eight characters in length; it is aligned under the MAJOR heading. The minor name, which corresponds to the *rname* field in the RESERVE macro instruction, can be from 1 to 255 characters in length. However, only 32 characters can appear in the report. When a minor name exceeds 32 characters, it is truncated in the report, and an asterisk (*) indicates the name has been truncated. If the minor name contains unprintable characters, RMF reports it in the form 'name', where *name* is shown in hexadecimal digits and is only 29 digits in length. Each minor name is aligned under the MINOR heading. RMF recognizes only 44 characters. Therefore if two minor names (both with the same major name) are longer than 44 characters and differ only beyond the forty- fourth character, RMF cannot distinguish between them.

Figure 4-40. Fields in the Reserve Activity Report

Transaction Activity Report

The transaction activity report enables you to evaluate the effectiveness of the performance objectives in the IPS and the workload classification in the installation control specification. The report is a "snapshot" of transaction activity at the time the report was requested. It shows, by performance group period, the transaction rate (the number of ended transactions per second) and the average elapsed time of each ended transaction.

You can compare the transaction rate and the average elapsed time required for each transaction to the objectives for the performance group specified in the IPS. When the installation control specification associates a performance group with a subsystem, a transaction class, a user identifier, or a transaction name, the report contains this data to enable you to relate the contents of the report to the work classification established in the installation control specification.

The data in the report is collected by the SRM and passed to RMF when a Monitor I session that is measuring workload activity is active. Thus, the transaction activity report can be obtained only when a Monitor I session with the WKLD option is running at the same time as the report request.

Figure 4-41 shows a sample transaction activity report. You can request transaction data for one or more individual performance groups or for all performance groups associated with a subsystem name. (For information on how to request the report, see "Menu Items" for a display session and "Monitor II Background Session Options" for a background session earlier in this chapter or "REPORTS Control Statement" for the post processor in Chapter 6.) Only performance groups with active work are included in the report; any performance group that had no activity does not appear. Any performance group specified in the report request that is not defined in the IPS or the installation control specification appears in the report marked as *****PGN NOT DEFINED*****.

RMF requires two "snapshots" in order to calculate the transaction rate and the average time for an ended transaction. Thus, TRANS RATE and AVG TRANS TIME contain dashes the first time the report is requested and the first time a performance group or performance group period appears in a series of report requests.

The field headings, their meanings, and the possible contents are described in Figure 4-42. Preceding the headings for the report are the time when the report was requested and the names of the installation control specification and the IPS in effect when the report was requested. The time is in the format hh:mm:ss, where hh is the hour, mm is the minute, and ss is the second, based on a 24-hour clock. The name of the installation control specification is in the format ICS=**name**, where **name** is the eight-character name of the IEAICSxx parmlib member or NONE if no IEAICSxx member is currently in effect. The name of the IPS is in the format IPS=**name**, where **name** is the eight-character name of the IPS or NONE if the coded defaults for the IPS are currently in effect.

To request a transaction activity report during a display session, specify:

```
TRX  [ALLPGN                    ]
     [subsystem-name            ]
     [performance-group,...     ]
```

```
  F                                CPU= 54 UIC=255 PDT=   3 DPR=  4   TRX     T
09:03:56  ICS=IEAICSAQ  IPS=IEAIPSAQ     ACCT     PERF    TRANS  AVG TRANS TIME
SUBSYS  TRXCLASS   USERID      TRXNAME     INFO   GRP PER   RATE    HHH.MM.SS.TTT
                                              C   0   1 -------  -------------
JES2                                       NO C   1   1 -------  -------------
JES2                                       NO         2 -------  -------------
TSO                                        NO C   2   1 -------  -------------
TSO                                        NO         2 -------  -------------
TSO                                        NO         3 -------  -------------
TSO                                        NO         4 -------  -------------
STC                                        NO C   3   1 -------  -------------
STC                            JES2        NO C   4   1 -------  -------------
STC                                        NO C   5   1 -------  -------------
JES2                           #(1)        NO C   6   1 -------  -------------
JES2                           #(1)        NO         2 -------  -------------
                                              C   7   1 -------  -------------
STC                            RMFGAT      NO C   8   1 -------  -------------
STC                            SONORA12    NO C   9   1 -------  -------------
TSO                D24RWP1                 NO C  10   1 -------  -------------
TSO                D24RWP1                 NO         2 -------  -------------
TSO                D24RWP1                 NO         3 -------  -------------
TSO                D24RWP1                 NO         4 -------  -------------
STC                            DSNMSTR     NO C  11   1 -------  -------------
STC                            IRLMPROC    NO C  12   1 -------  -------------
```

Figure 4-41. Transaction Activity (TRX) Report

Field Heading	Meaning
SUBSYS	One to four character name of the subsystem associated with the performance group in the installation control specification. When a subsystem is specified by name in the report request, all performance groups associated with that subsystem are grouped in the report. The field is blank when there is no installation control specification, when it does not define a subsystem name for the performance group, or when the name is not unique (that is, when more than one subsystem name is associated with the performance group).
TRXCLASS	One to ten character name of the transaction class defined in the installation control specification for the performance group. If a substring was used to define the name, the name reported includes the substring position number. The field is blank when there is no installation control specification, when it does not define a transaction class for the performance group, or when the name is not unique (that is, when more than one transaction class name is associated with the performance group).
USERID	One to ten character user identifier defined in the installation control specification for the performance group. If a substring was used to define the user identifier, the user identifier reported includes the substring position number. The field is blank when there is no installation control specification, when it does not define a user identifier for the performance group, or when the user identifier is not unique (that is, when more than one user identifier is associated with the performance group).
TRXNAME	One to ten character transaction name defined in the installation control specification for the performance group. If a substring was used to define the same, the name reported includes the substring position number. The field is blank when there is no installation control specification, when it does not define a transaction name for the performance group, or when the transaction name is not unique (that is, when more than one transaction name is associated with the performance group).
ACCT INFO	'YES' indicates that an account number was specified in the installation control specification for the performance group. Blanks indicate that no account number has been assigned.
PERF GRP	One to three digit number in the range of 0 through 999 that identifies the performance group. The number is preceded by a C when the performance group is a control performance group and by an R when the performance group is a report performance group.
PERF PER	One digit number in the range of 1 through 8 that identifies the performance group period. When more than one period is active for a performance group, this field, TRANS RATE, and AVG TRANS TIME are repeated for each active performance group period.
TRANS RATE	Number in the form nnn.nnn that indicates the transaction rate (number of ended transactions per second) for the performance group period. Because RMF requires two "snapshots" to compute the value for the field, TRANS RATE contains dashes the first time the report is requested or the first time a performance group period appears in a series of reports. TRANS RATE is repeated for each performance group period included for a performance group. To calculate the transaction rate RMF uses the time that has elapsed between two requests for the report and the number of transactions that ended during the elapsed time. RMF divides the number of ended transactions by the length of the elapsed time to determine the transaction rate.

Figure 4-42 (Part 1 of 2). Fields in the Transaction Activity Report

Field Heading	Meaning
AVG TRANS TIME HHH.MM.SS.TTT	The average elapsed time of all ended transactions for the performance group period, where HHH is the hours, MM is the minutes, SS, is the seconds, and TTT is thousandths of a second. Because RMF requires two "snapshots" to compute the value for the field, AVG TRANS TIME contains dashes the first time the report is requested or the first time a performance group appears in a series of reports. AVG TRANS TIME is repeated for each performance group period includes for a performance group. To calculate the average transaction time, RMF determines the total number of transactions that ended during the time between the report requests and computes the total elapsed time for all ended transactions. RMF then divides the total time for all ended transactions by the total number of ended transactions to determine the average transaction time.

Figure 4-42 (Part 2 of 2). Fields in the Transaction Activity Report

Chapter 5. Post Processor

The RMF post processor can generate interval, duration, summary, exception and plot reports, thus enabling your installation to obtain interval report during off-shift hours instead of during RMF processing and to obtain reports that summarize the information collected by RMF during a Monitor I session.

There are two types of input to the post processor: data input and control input.

Data Input

Data input consists of a data set that contains SMF records. The data sets that SMF produces are VSAM data sets named SYS1.MANn, where **n** is an alphabetic character (A-Z) or a numeric character (0-9). You must dump the records in an SMF data set by means of the IFASMFDP program to a non-VSAM data set because the post processor cannot accept input from a VSAM data set. Your data set should include only the records that RMF (types 70-79) produces because the post processor cannot accept records greater than 32756 bytes. For more information on the IFASMFDP program, see *System Management Facilities (SMF)*.

The input stream can contain records archived from RMF Version 2. For a description on processing archived records, see Chapter 8.

The input records **must** be in order by interval start date and interval start time. If your input data set combines records from several data sets, the records must be dumped in the order in which they were written or sorted before they are used as input to the post processor. If your input data set combines records created by more than one system, sort the records into one data stream. If you sort records, you must preserve the original order of the records, within interval date and time.

Use the SYSID control statement, described later in this chapter, to tell the post processor whether you want a report for a specific system or separate reports for all systems included in the input data set.

Because the interval start date and time fields are not always in the same location for each record, use ERBPPSRT, a sort exit supplied with RMF, to resolve the date and time fields. The following sample uses the sort provided in OS/VS SORT/Merge Program Product 5740-SM1. So the sort program can sort by move-to-buffer date and time, use ERBPPSRT as the input (E15) and output (E35) exits. ERBPPSRT exchanges the interval start date and time with move-to-buffer date and time. After sorting, ERBPPSRT exchanges the fields to give the records their original contents. The sort exit can handle records archived from RMF Version 2 as well as those generated by RMF Version 3.

```
//SAMPLE   JOB   accounting information
//SORT     EXEC  SORTD
//SORTIN   DD    DSN=D09BAB1.RMFV7.DATA,DISP=SHR
//         DD    DSN=D09BAB1.RMFV5.DATA,DISP=SHR
//SORTOUT  DD    DSN=&SMFOUT,DISP=(,PASS),UNIT=SYSVIO,
//               SPACE=(TRK,(200,200))
//SORTWK01 DD    SPACE=(TRK,200),UNIT=SYSVIO
//SORTWK02 DD    SPACE=(TRK,200),UNIT=SYSVIO
//SORTWK03 DD    SPACE=(TRK,200),UNIT=SYSVIO
//EXITLIB  DD    DSN=SYS1.LINKLIB,DISP=SHR
//SYSIN    DD    *
    SORT FIELDS=(11,4,CH,A,7,4,CH,A),EQUALS
    MODS E15=(ERBPPSRT,500,EXITLIB,N),
             E35=(ERBPPSRT,500,EXITLIB,N)
/*
//ALLSYSID EXEC  PGM=ERBRMFPP
//MFPINPUT DD    DSN=&SMFOUT,DISP=(OLD,PASS)
//MFPMSGDS DD    SYSOUT=*
//SYSIN    DD    *
REPORTS (CPU)
SYSOUT (A)
/*
```

For more information on SORTD, E15 and E35, see *OS/VS SORT/Merge Programmer's Guide* SC33-4035.

Control Input

Control input consists of control statements. The control statements describe the type of reporting you want the post processor to do.

Output from the post processor consists of the printed reports you requested and any messages issued by the post processor.

The post processor verifies the control statements and uses the information they contain to build a list of options that control its execution. If you omit a statement, RMF substitutes the default value. If a syntax error is detected, RMF substitutes the default value, if there is one, or ignores the option. In most cases, an error in a control statement does not cause the post processor to terminate; it notes the condition, issues a warning message to the output message data set, and continues with the process of building a list of options for execution. After this process is complete, the post processor issues a message to the output message data set to indicate the options in effect.

The option list consists of the options you have entered on control statements and any options for which the default values were used.

Each option listed is followed by the input source from which the post processor obtained the option. The possible sources are:

- SYSIN -- the option was specified on a control statement for the post processor.
- DEFAULT -- the option was taken from the control statement defaults. When RMF detects an invalid value and substitutes a default value, a warning message is issued, and DEFAULT appears in the option list.

After establishing the list of options, the post processor generates the reports you specify for the reporting period you specify. The duration of the reporting period can be established in two ways:

1. Through the use of control statements (DATE, RTOD, PTOD, ETOD, and STOD) that indicate a specific range of dates and specific ranges of times.
2. Through the contents of the input record data set. The defaults for the control statements that define the reporting period cause the post processor to include all dates and all times in the reporting period. Thus, omitting these control statements causes the requested reports to be generated for all dates and times included in the input data set.

A single execution of the post processor can generate plot reports, exception reports, exception-generated interval reports, summary reports, and either interval or duration reports. The use of control statements to describe the type of reporting you require is given later in this chapter under "Post Processor Control Statements" on page 5-5.

When all requested reports have been generated, the post processor issues a return code and terminates execution. The return codes from the post processor are:

Code	Meaning
0	Normal completion -- reports generated as requested
4	Normal completion -- no RMF input records found or no RMF input records found that meet the user requirements specified in the control statements
8	Error -- see accompanying RMF message
12	Terminating error -- see accompanying RMF message

Post Processor JCL

The post processor runs as a batch job; the JCL required to invoke the post processor is:

```
//EXAMPLE  JOB  accounting information
//POST     EXEC PGM=ERBRMFPP
//MFPINPUT DD   data set contains SMF records (tape or DASD)
//SYSIN    DD     *
  control statements
/*
```

To run the RMF post processor while the system is operating in 370 mode, you must add a JOBLIB or STEPLIB DD statement to the previous JCL. This DD statement identifies SYS1.LINKLIB in your MVS/XA system. The following example shows the JCL statements needed to run the post processor in 370 mode:

```
//EXAMPLE  JOB  accounting information
//POST     EXEC PGM=ERBRMFPP
//STEP=IB  DD   DSN=SYS1.LINKLIB,DISP=SHR,
//              VOL=SER=volume containing MVS/XA linklib,
//              UNIT=unit type
```

Because RMF can generate spanned SMF records -- particularly when I/O device activity is measured -- the DD statement for the MFPINPUT data set must contain DCB parameters if the input is an unlabeled tape. Specify the following DCB parameters on the DD statement: RECFM=VBS,LRECL=32756, BLKSIZE=xxx. If the input is a labelled tape or DASD, do not specify any DCB parameters.

The control statements can appear in any order. Data can appear in columns 1 to 72. No continuation statements are allowed; however, a control statement can be repeated until all required options are specified. The maximum number of control statements that can be specified in a single execution of the post processor is 716. Each control statement is described in detail in the following section. Examples of using control statements to describe the types of reporting required are given under "Examples of Control Statements" on page 5-36.

The post processor dynamically allocates all message and report data sets to SYSOUT, unless the requested DDNAME has been preallocated. However, if the system returns a code that indicates the DDNAME is in use, the post processor proceeds as if the dynamic allocation had been successful. Thus, your installation can allocate the data sets for the post processor in the JCL, enabling you to change some characteristics of the output. The report data sets for Monitor I interval reports and duration reports use the DDNAMEs MFRnnnnn and MFEnnnnn, which are described in Chapter 2 under "RMF JCL." The other report data sets and the message data set that the post processor uses and the DDNAMEs for these data sets are:

DDNAME	Use
MFPMSGDS	This data set contains message output. One MFPMSGDS data set is allocated each time the post processor is executed. **Note:** The SYSOUT class for this data set cannot be changed by the SYSOUT (class) parameter in the control statements; to change the SYSOUT class, you must preallocate the data set.
MFXnnnnn	This data set contains exception interval report data. A DDNAME is generated and a data set created for each interval included in the exception report; nnnnn is a decimal number from 00001 to 99999. The first interval is assigned the DDNAME MFX00001. The second MFX00002, and each subsequent interval is assigned a DDNAME in ascending numerical order. If no exception interval reports are produced in a given interval, the data set for that interval is empty.
RMFPnnnn	This data set contains Monitor II session interval report output. A DDNAME is generated and a data set is created for each report for each session-identifier included in the reporting; **nnnn** is a decimal number from 0001 to 9999. When the post processor is to generate reports for more than one system, a separate data set is allocated for each report for each system. When operands for a Monitor II session were not specified on the REPORTS statement, the post processor uses the operands in the SMF record, and a separate data set is allocated each time the operands change.
PPSUMnnn	This data set contains summary report output. A DDNAME is generated and a data set is created for each system included in the input data set; **nnn** is a decimal number from 001 to 999. The first system encountered is assigned the DDNAME PPSUM001, and each subsequent system is assigned a DDNAME in ascending numerical order.
PPPLTnnn	This data set contains plot report output. A DDNAME is generated and a data set is created for each system included in the input data set; **nnn** is a decimal number from 001 to 999. The first system encountered is assigned the DDNAME PPPLT001, and each subsequent system is assigned a DDNAME in ascending numerical order. The data set the post processor creates has a variable blocked (VBA) record format and a logical record length of 137, a length that allows for two plot reports per page. If your printer has a longer line length, you can pre-allocate this data set with a logical record length of 193 to allow three plots per page. A logical record length of 193 allows 186 characters for the three 62-character plots, two separator characters between the plots, one ASA control character, and four characters for the variable record header.
PPEXTnnn	This data set contains exception report output. A DDNAME is generated and a data set is created for each system included in the input data set; nnn is a decimal number from 001 to 999. The first system encountered is assigned to DDNAME PPEXT001, and each subsequent system is assigned a DDNAME in ascending numerical order.

Note: When you omit the DCB characteristics for the message and report data sets described above, the characteristics used are:

```
DCB=(RECFM=VBA,LRECL=137,BLKSIZE=1693)
```


When you change the DCB characteristics, you cannot change the record format; you must specify RECFM=VBA.

Post Processor Control Statements

This section describes in alphabetical order each of the post processor control statements. Figure 5-1 shows which control statements are needed for interval, summary, duration, exception, and plot reports. It also indicates which control statements are required and which control statements can be omitted when the default value is what you want.

Post	Used For:						
Processor Control Statement	Monitor I Session Interval Report	Monitor II Session Interval Report	Duration Report	Summary Report	Plot Report	Exception Report	Notes
DATE	X	X	X	X	X	X	1
DELTA		X					1
DINTV			X				2,3
ETOD						X	1
EXCEPT						X	2
EXITS	X	X	X	X	X	X	1
EXRPTS						X	2
MAXPLEN					X		1
PINTV					X		1
PLOTS					X		2
PTOD					X		1
REPORTS	X	X	X				2
RTOD	X	X	X				1
SESSION		X					1
STOD				X			1
SUMMARY				X			1
SYSID	X	X	X	X	X	X	1
SYSOUT	X	X	X	X	X	X	1

Notes:
1. The control statement need not be explicitly specified if the default value is acceptable.
2. The control statement must be specified explicitly.
3. Duration reports cannot be requested concurrently with interval reports; separate executions of the post processor are required for each type. However, duration reports can be requested concurrently with exception generated interval reports; summary, exception, and plot reports can be requested concurrently with either duration or interval reports.

Figure 5-1. Post Processor Control Statements Summary

DATE Control Statement

The DATE control statement specifies the date or range of dates of the reporting period for all reports.

The syntax of the statement is:

```
DATE(yyddd,yyddd)
```

Where **yy** is the last two digits of the year, and **ddd** is the day of the year. The first date is the starting date, and the second date is the ending date. The dates must be specified in chronological order.

When the DATE statement is omitted, the default value is DATE(00000,99366); that is, all dates found in the input data set are reported.

When the entire reporting period falls within one calendar day, the date must be specified twice. For example, to request reports for data collected on June 7, 1982, specify DATE(82158,82158).

DELTA Control Statement

The DELTA control statement specifies whether certain fields in Monitor II background session reports are to reflect total values or, after the first report, changed -- "delta" -- values. The fields that are affected by delta mode are described for each report under "Monitor II Session Reports" in Chapter 4. The syntax of the statement is:

```
DELTA/NODELTA
```

DELTA indicates that the affected fields are to reflect changed values; that is, the reports are to be generated in delta mode. NODELTA indicates that the affected fields are to reflect total values. NODELTA is the default.

DINTV Control Statement

The DINTV control statement specifies that the post processor is to generate duration reports and indicates the length of the duration interval. The duration interval is the length of time each report can cover. The syntax of the statement is:

```
DINTV(hhmm)
```

Where **hh** is the hour and **mm** is the minute on a 24-hour clock. The time specified is the length of the duration interval, not the time of day when the interval is to begin.

The DINTV statement must be specified to obtain duration reports. When the statement is omitted, the post processor generates interval reports. When the statement is specified, the length of the duration interval must be specified because there is no default value for the length of the duration interval.

The duration interval can be the same length as the reporting period. It can also be less than the reporting period; in this case, the reporting period consists of several duration intervals. Figure 5-2 illustrates how the duration interval relates to the reporting period. Assume a reporting period that covers the twelve-hour block of time from 8:00 A.M. to 8:00 P.M. over a range of three days. As the figure shows, specifying DINTV(1200) causes the post processor to generate three duration reports, each covering twelve hours of system activity. Specifying DINTV(0600) causes the post processor to generate six duration reports, each covering six hours of system activity. You could also choose a duration interval that does not break the reporting period into equal blocks of time. As the figure shows, specifying DINTV(1000), indicating a duration interval of ten hours, would cause the post processor to generate the following reports:

1. 8 A.M. to 6 P.M. on the first day of the reporting period.
2. 6 P.M. to 8 P.M. on the first day of the reporting period, and 8 A.M. to 4 P.M. on the second day.

3. 4 P.M. to 8 P.M. on the second day of the reporting period, and 8 A.M. to 2 P.M. on the third day.
4. 2 P.M. to 8 P.M. on the third day of the reporting period.

In this case, if you wanted to use the reports to compare system performance over the same hours on each day of the reporting period, it would be difficult because each report covers a different time range and some span two days and include a twelve-hour gap when no reporting was done.

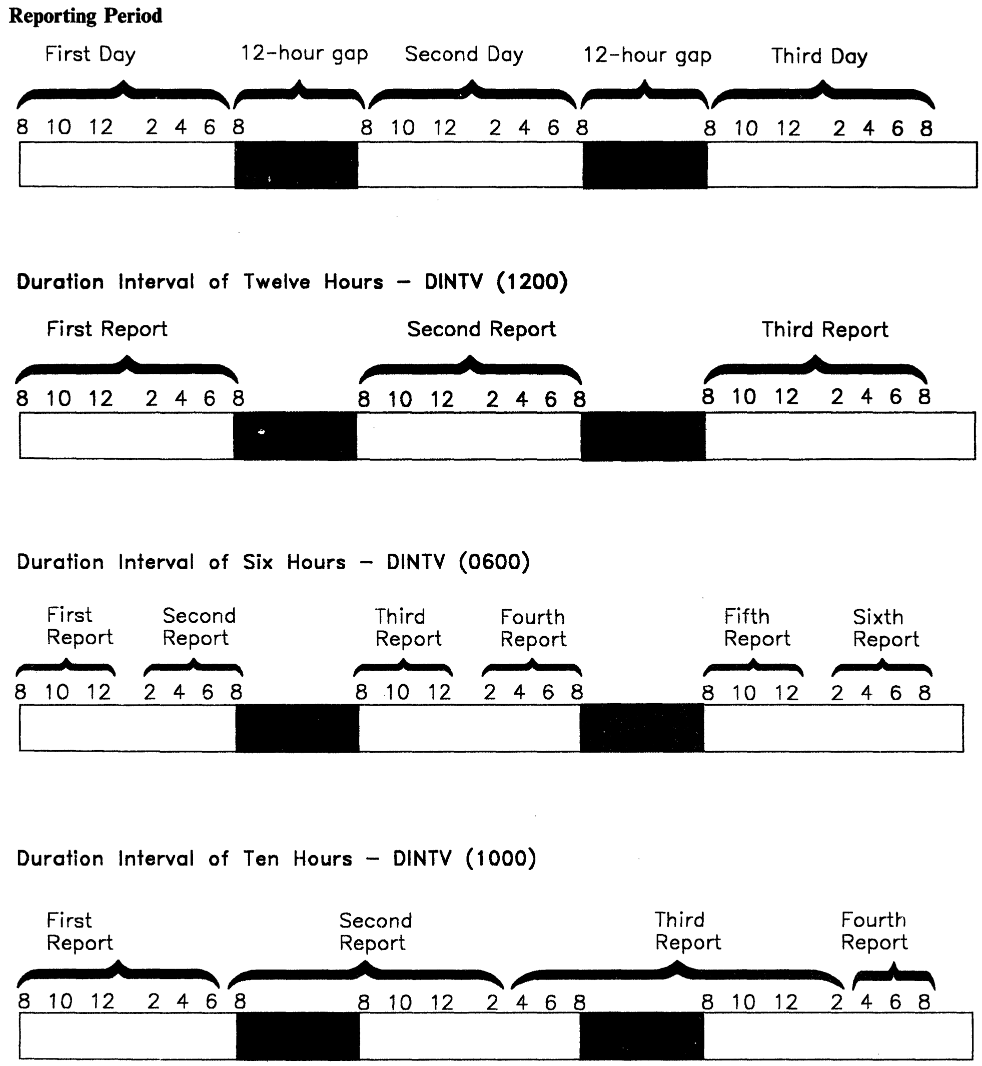

Figure 5-2. Relationship Between Duration Interval and Reporting Period

The syntax of the control statement allows a maximum duration interval of 99 hours and 60 minutes; however, the recommended maximum length is 24 hours. A duration interval greater than 24 hours could cause data values for report fields to overflow the space allowed in the formatted report. Such an overflow condition is indicated by a negative number or asterisks in the field, an indication that the length of the duration interval should be reduced.

Thus, care should be taken in selecting a duration interval, and, for most effective reporting, the duration interval should bear a relationship to the length of the reporting period.

The actual length of time included in the resulting duration report depends on the time within the duration interval when data was actually collected. For example, the following post processor control statements request a duration report for channel path activity that has a duration interval of 12 hours (from 8:00 A.M. to 8:00 P.M.):

```
RTOD(0800,2000)
DINTV(1200)
REPORTS(CHAN)
```

However, if channel path activity measurement did not begin until 10:00 A.M., and ended at 6:00 P.M., the duration report covers the time from 10:00 A.M. to 6:00 P.M. The start time of the first and last RMF measurement intervals actually included in the report appear in the heading for the report.

Note: Because the length of the RMF measurement intervals can vary by a few seconds or more, the number of intervals that the post processor includes in a duration interval can also vary.

ETOD Control Statement

The ETOD control statement specifies the starting time and ending time of the reporting period for an exception report for each day in the reporting period.

The syntax of the statement is:

```
ETOD(hhmm,hhmm)
```

Where **hh** is the hour and mm is the minute on a 24-hour clock. The first time specifies the beginning of the reporting period, and the second time specifies the end of the reporting period. The second must be later than the first, or a syntax error occurs.

When the ETOD statement is omitted, the default value is ETOD(0000,2400); that is, all times are reported. Thus, you would use this control statement when you want a reporting period for an exception report that is different from the default value.

Because the range of values allowed is from 0000 to 2400, it is not possible to define a reporting period that consists of a single block of time that spans more than one calendar day. Thus, you cannot define a reporting period that, for example, runs from 12 noon on one day to 12 noon on the next day. However, you can define a reporting period that consists of the same block of time over several days. For example, to produce an exception report using data collected from 8:00 A.M. to 1:00 P.M. for the week beginning on January 3, 1982 and ending on January 9, 1982, the required DATE and ETOD statements would be:

```
DATE(82003,82009)
ETOD(0800,1300)
```


EXCEPT Control Statement

The EXCEPT control statement defines a condition that RMF is to test for an exceptional value. The definition consists of a condition – a system indicator that RMF recognizes by name – a threshold value for the condition, and an operator that establishes the relation between the condition and the threshold value. RMF compares the threshold value for the condition with the contents of the appropriate SMF record field; if the condition exceeds the threshold, RMF recognizes that an exception has occurred. A table of the conditions that RMF can recognize and test for exceptional values follows this section. An exception consists of one condition or several conditions; when you group conditions into a single condition, all conditions must exceed their threshold values in order for RMF to recognize that an exception has occurred.

When an exception occurs, RMF writes a line to the exception report to describe the exception. By supplying an EXRPTS control statement for the exception, you can also cause RMF to generate one or more interval reports when the exception occurs.

Only one condition and threshold can be specified in a single control statement, but you specify as many EXCEPT statements as you require. The syntax of the EXCEPT statement is:

```
                                                                   [,LE ]
EXCEPT([control-statement-name](condition-name[(qualifier)] [,GE ],threshold-value))
```

control-statement-name
Specifies a one to eight character name, starting with an alphabetic character, that has two uses. First, it provides a means of grouping multiple conditions to form a single exception. You group conditions by coding the same control-statement-name on each separate EXCEPT statement that defines each one of the conditions that form the exception. When conditions are grouped, all conditions must be satisfied in order for RMF to recognize the exception.

Second, the control-statement-name associates the exception with the interval reports, if any, that RMF is to generate when the exception occurs. The EXRPTS statement defines these reports. For more information, see the EXRPTS control statement.

When exception reports are generated, exceptions are listed in chronological order by control-statement-name. The control-statement-names are listed in alphabetical order. Choosing a meaningful control-statement-name makes it easier to recognize an exception.

condition-name
Specifies the name of the condition that RMF is to test for an exceptional value. RMF performs the test by comparing the contents of the appropriate field in an SMF record with the threshold value. The conditions that can be specified are listed in the tables following this section. The tables list the conditions according to the record type that contains the data that RMF compares with the threshold. Thus, RMF can recognize an exception only when the required SMF record was produced during the reporting period.

qualifier
Specifies an optional qualifier that can limit the scope of the condition identified by condition-name. The following tables also list the qualifiers.

LE or GE

Specifies the relational operator RMF is to use to determine if the exception has occurred.

LE indicates that any value in the SMF record that is less than or equal to the threshold value causes an exception.

GE indicates that any value in the SMF record that is greater than or equal to the threshold value causes an exception.

threshold-value

Specifies the value that RMF compares to a computed value from the appropriate SMF record fields. The value can be specified as:

- a whole number – where the value can be a one to six digit integer.
- a fraction – up to six digits can appear before and up to three digits can appear after the decimal point. However, the total number of digits specified cannot exceed nine, including the decimal point.
- percentage – where the maximum percentage that can be specified is 100 percent. The value can be expressed as a whole integer or as a fraction. You can specify a percentage only with those conditions that indicate a percent value.

  When the SMF record exceeds the threshold value, as indicated by the LE or GE operand, RMF flags the condition for reporting. For each interval, only one line is printed for each exception regardless of the number of times the threshold is exceeded.

Examples

1. The post processor is to generate a line in the exception report for each interval when the percent busy value for channel path 1 is greater than or equal to ten percent. Use the following control statement:

   ```
   EXCEPT((CHPBSY(01),GE,10))
   ```

   The absence of a control-statement-name indicates that this is a single condition exception and that no interval reports are generated.

2. The post processor is to generate a line in the exception report for each interval when the I/O service rate for performance group number three is less than or equal to .01 and the device percent utilized value for device 6D8 is greater than or equal to 3. Use the following control statements:

   ```
   EXCEPT(EXCP1(IOSRV(0030),LE,.01))
   EXCEPT(EXCP1(DVUTL(6D8),GE,3))
   ```

   Note that the control-statement-name, EXCP1, is used to group the two conditions into one exception. Only when both conditions are recognized will RMF write a line for each condition in the exception report.

Exceptions based on Type 70 (CPU Activity) SMF record				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
percent processor busy	CPUBSY	cpuid[1]	SMF70WAT SMF70CID SMF70INT SMF70INT SMF70CPN	If a processor identifier is specified, RMF: 1. Finds the processor section for the specified processor. 2. Calculates the percent busy from the processor wait time. 3. Compares the percent busy against the exception threshold. If a processor identifier is not specified, RMF: 1. Examines all processor data sections. 2. Calculates the percent busy for each processor from the processor wait time. 3. Calculates the average percent busy for all processors. 4. Compares the average percent busy against the threshold value.
maximum number of batch users	MXBATCH	none	SMF70BMM	direct comparison
maximum number of started tasks	MXSTC	none	SMF70SMM	direct comparison
maximum number of TSO users	MXTSO	none	SMF70TMM	direct comparison
average number of batch jobs	AVGBATCH	none	SMF70BTT SMF70SAM	BTT/SAM
average number of started tasks	AVGSTC	none	SMF70STT SMF70SAM	STT/SAM
average number of TSO users	AVGTSO	none	SMF70TTT SMF70SAM	TTT/SAM
average number of in and ready users	AVGIARDY	none	SMF70RTT SMF70SAM	RTT/SAM
average number of out and ready users	AVGOARDY	none	SMF700TT SMF70SAM	0TT/SAM

[1]cpuid specifies a one-digit processor identifier
If the qualifier is omitted, the threshold applies to the average of all processors.

Figure 5-3. Exception Control Statements based on Type 70 (CPU Activity) SMF record

Exceptions based on Type 71 (Paging Activity) SMF record				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
size of real storage (K)	STORAGE	none	SMF71TFC SMF71FIN	TFC+FIN
total number of pages per second	TPAGRT	none	SMF71PIN SMF71POT SMF71SIN SMF71SOT SMF71VIN SMF71VOT SMF71SNI SMF71SNO SMF71INT	(PIN+POT+SIN+SOT+ VIN+VOT+SNI+SNO)/INT
swap out due to detected long think time[2]	PSOCLTT	none	SMF71LFS SMF71INT	LFS/INT
swap out due to long wait time[3]	PSOCLW	none	SMF71TOT(3) SMF71INT	TOT(3)/INT
swap out due to long wait time[4]	PSOCLW	none	SMF71AXD(3) SMF71LES(3) SMF71LAX(3) SMF71ESD(3) SMF71MIG(3) SMF71INT	(AXD(3))+LES(3)+ LAX(3)+ESD(3)+MIG(3))/INT
swap-out due to detected wait[3]	PSOCDW	none	SMF71TOT(6) SMF71INT	TOT(6)/INT
swap-out due to detected wait[4]	PSOCDW	none	SMF71AXD(6) SMF71LES(6) SMF71LAX(6) SMF71ESD(6) SMF71MIG(6) SMF71INT	(AXD(6)+LES(6)+ LAX(6)+ESD(6)+MIG(6))/INT
number of unilateral swap-outs[3]	PSOCU	none	SMF71TOT(10) SMF71INT	TOT(10)/INT
number of unilateral swap-outs[4]	PSOCU	none	SMF71AXD(10) SMF71LES(10) SMF71LAX(10) SMF71ESD(10) SMF71MIG(10) SMF71INT	(AXD(10)+LES(10)+ LAX(10)+ESD(10)+MIG(10))/INT
swap-out on recommendation value[3]	PSOCEORV	none	SMF71TOT(9) SMF71INT	TOT(9)/INT
swap-out on recommendation value[4]	PSOCEORV	none	SMF71AXD(9) SMF71LES(9) SMF71LAX(9) SMF71ESD(9) SMF71MIG(9) SMF71INT	(AXD(9)+LES(9)+ LAX(9)+ESD(9)+MIG(9))/INT
swap-out on enqueue exchange[3]	PSOCENQE	none	SMF71TOT(8) SMF71INT	TOT(8)/INT

[2] This option is no longer supported for records created by RMF Version 3 Release 3.
[3] for SMF records created by RMF releases before Version 3 Release 3
[4] for SMF records created by RMF Version 3 Release 3

Figure 5-4 (Part 1 of 3). Exception Control Statements based on Type 71 (Paging Activity) SMF record

Exceptions based on Type 71 (Paging Activity) SMF record				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
swap-out on enqueue exchange[4]	PSOCENQE	none	SMF71AXD(8) SMF71LES(8) SMF71LAX(8) SMF71ESD(8) SMF71MIG(8) SMF71INT	(AXD(8) + LES(8) + LAX(8) + ESD(8) + MIG(8))/INT
number of requested swap-outs[3]	PSOCREQ	none	SMF71TOT(7) SMF71INT	TOT(7)/INT
number of requested swap-outs[4]	PSOCREQ	none	SMF71AXD(7) SMF71LES(7) SMF71LAX(7) SMF71ESD(7) SMF71MIG(7) SMF71INT	(AXD(7) + LES(7) + LAX(7) + ESD(7) + MIG(7))/INT
swap-out due to auxiliary storage shortage[3]	PSOCAXSS	none	SMF71TOT(4) SMF71INT	TOT(4)/INT
swap-out due to auxiliary storage shortage[4]	PSOCAXSS	none	SMF71AXD(4) SMF71LES(4) SMF71LAX(4) SMF71ESD(4) SMF71MIG(4) SMF71INT	(AXD(4) + LES(4) + LAX(4) + ESD(4) + MIG(4))/INT
swap-out due to transition to non-swappable[3]	PSOCTONS	none	SMF71TOT(11) SMF71INT	TOT(11)/INT
swap-out due to transition to non-swappable[4]	PSOCTONS	none	SMF71AXD(11) SMF71LES(11) SMF71LAX(11) SMF71ESD(11) SMF71MIG(11) SMF71INT	(AXD(11) + LES(11) + LAX(11) + ESD(11) + MIG(11))/INT
number of page faults per second	PAGERT	none	SMF71PIN SMF71PRC SMF71INT	(PIN + PRC)/INT
demand paging per second	DPAGRT	none	SMF71PIN SMF71POT SMF71INT	(PIN + POT)/INT
swap rate	SWART	none	SMF71SSQ SMF71INT	SSQ/INT
percent successful swap-in[2]	PLSWAPIN	none	SMF71LSS SMF71LGS SMF71LFS SMF71LSE	((LSS + LGS-LFS-LSE)*100)/LGS

[2] This option is no longer supported for records created by RMF Version 3 Release 3.
[3] for SMF records created by RMF releases before Version 3 Release 3
[4] for SMF records created by RMF Version 3 Release 3
[5] K indicates that the values are for all swap reasons.

Figure 5-4 (Part 2 of 3). Exception Control Statements based on Type 71 (Paging Activity) SMF record

Exceptions based on Type 71 (Paging Activity) SMF record				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
percent successful swap-out[3]	PLSWAPOU	none	SMF71LGS SMF71TOT(2) SMF71TOT	(LGS*100)/(TOT(2)+TOT)
percent successful swap-out[4]	PLSWAPOU	none	SMF71TOT(k) SMF71AXD(k) SMF71ESD(k) SMF71LES(k) SMF71LAX(k)	(TOT(k)-AXD(k)-ESD(k) -LES(k)-LAX(k))*100/(TOT(k)-AXD(k) -ESD(K))[5]
maximum number of SQA frames	MXSQA	none	SMF71MXQ	direct comparison
average number of SQA frames	AVGSQA	none	SMF71AVQ	direct comparison
maximum number of CSA fixed-frames	MXCSAF	none	SMF71MXC	direct comparison
maximum number of total CSA frames	MXCSAT	none	SMF71MXP	direct comparison
average number of total CSA frames	AVGCSAT	none	SMF71AVP	direct comparison
average number of CSA fixed frames	AVGCSAF	none	SMF71AVC	direct comparison
average number of VIO allocated slots	AVGVIOF	none	SMF71AVV	direct comparison
maximum number of allocated VIO slots	MAXVIOF	none	SMF71MXV	direct comparison
swap-out terminal wait[3]	PSOCTW	none	SMF71TOT(2) SMF71TOT(1)	(TOT(2)+TOT(1)-LGS)/INT
swap-out terminal wait[4]	PSOCTW	none	SMF71AXD(1) SMF71LES(1) SMF71LAX(1) SMF71ESD(1) SMF71MIG(1) SMF71AXD(2) SMF71LES(2) SMF71LAX(2) SMF71ESD(2) SMF71MIG(2) SMF71INT	((AXD(1)+LES(1)+LAX(1)+ESD(1)+ MIG(1)+AXD(2)+LES(2)+LAX(2)+ ESD(2)+MIG(2))/INT
page move rate	PGMVRT	none	SMF71PMV SMF71INT	PMV/INT
swap-out due to real storage shortage	PSOCRPSS	none	SMF71RSS SMF71INT	RSS/INT
average high unreferenced interval count for real storage frames[6]	AVGHUIC	none	SMF71ACA	ACA*100
maximum high unreferenced interval count for real storage frames[6]	MXHUIC	none	SMF71HIC	HIC*1000
page movement rate to extended storage[6]	PTES	none	SMF71PES	PES/INT
migration rate from extended to auxiliary storage[6]	ESMR	none	SMF71PEA	PEA/INT

[3] for SMF records created by RMF releases before Version 3 Release 3
[4] for SMF records created by RMF Version 3 Release 3
[6] Exceptions for these SMF records are calculated only if RMF is running on a 3090 processor.

Figure 5-4 (Part 3 of 3). Exception Control Statements based on Type 71 (Paging Activity) SMF record

 LC28-1556-2

Exceptions based on Type 72 (Workload Activity) SMF record				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
total service per second	TOTSRV	pgnp or subsys or dmn[7]	SMF72SER SMF72INT	SER/INT
I/O service rate per second	IOSRV	pgnp or subsys or dmn[7]	SMF72ITS SMF72INT	ITS/INT
processor service per second	CPUSRV	pgnp or subsys or dmn[7]	SMF72CTS SMF72INT	CTS/INT
total SRB service per second	SRBSRV	pgnp or subsys or dmn[7]	SMF72sts SMF72INT	STS/INT
main storage service per second	MSOSRV	pgnp or subsys or dmn[7]	SMF72MTS SMF72INT	MTS/INT
number of transactions per second	TRANS	pgnp or subsys or dmn[7]	SMF72TTX SMF72INT	TTX/INT
response time	RTIME	pgnp or subsys or dmn[7]	SMF72TTM SMF72TTX	TTM/TTX
number of swaps per transaction	SPERTRA	pgnp or subsys or dmn[7]	SMF72TTX SMF72SPP	TTX/SPP

[7]pgnp specifies a number in the form nnnp where nnn is a three digit performance group number in the range of 000 through 999, and p is a one digit performance group period in the range 1 through 8. When you specify a performance group period of zero, the threshold applies to the sum of all performance group periods of the performance group. When you specify a performance group period of 1 through 8, the threshold applies only to the specified performance group period for the performance group. If the qualifier is omitted, the threshold applies to the sum of all performance group periods.

subsys specifies a one to four-byte subsystem name; the first character must be alphabetic, and the last three can be alphabetic or numeric.

dmn specifies a one to three digit domain number in the range 0 through 128.

Figure 5-5. Exception Control Statements based on Type 72 (Workload Activity) SMF record

Exceptions based on Type 73 (Channel Path Activity) SMF record				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
percent channel busy	CHPBSY	cpid[8]	SMF73BSY SMF73SAM	(BSY*100)/SAM

[8]cpid specifies a two-digit hexadecimal number that identifies a channel path. Specifying the qualifier is optional; if it is omitted, the threshold applies to all channel paths in the SMF record.

Figure 5-6. Exception Control Statements based on Type 73 (Channel Path Activity) SMF record

Exceptions based on Type 74 (Device Activity) SMF record[9]				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
percent not ready	DNOTRDY	devnmbr, volser, stg grp, or class[10]	SMF74NRD SMF74SAM	(NRD*100)/SAM
percent reserved	DR	devnmbr, volser, stg grp, or class[10]	SMF74RSV SMF74SAM	(RSV*100)/SAM
percent mount pending	DMTPEND	devnmbr, volser, stg grp, or class[10]	SMF74MTP SMF74SAM	(MTP*100)/SAM
percent device utilization	DVUTL	devnmbr, volser, stg grp, or class[10]	SMF74CNN SMF74DIS SMF74INT SMF74UTL SMF74SAM	$100* \frac{(CNN+DIS)}{\text{INT in secs.}} + \frac{UTL}{SAM}$
device activity rate	DART	devnmbr, volser, stg grp, or class[10]	SMF74SSC SMF74INT	SSC/INT in seconds
average connect time	DCTAVG	devnmbr, volser, stg grp, or class[10] .	SMF74CNN SMF74MEC	CNN/MEC
average disconnect time	DDTAVG	devnmbr, volser, stg grp, or class[10]	SMF74DIS SMF74MEC	DIS/MEC
average pending time	DPTAVG	devnmbr, volser, stg grp, or class[10]	SMF74PEN SMF74MEC	PEN/MEC
average IOS queue time	DQTAVG	devnmbr, volser, stg grp, or class[10]	SMF74QUE SMF74SAM SMF74SSC SMF74INT	(QUE/SAM)/(SSC/INT)

[9] specifying a condition based on this record, you must use a qualifier. If a qualifier is not used, a syntax error occurs and the condition is not processed.

[10]devnmbr specifies a one to three-digit hexadecimal device number in the range 000 through FFF.
volser specifies a one to six character volume serial number enclosed in quotes.
stg grp specifies a one to eight character storage group name
class specifies any of the six valid device classes for Monitor I device activity measurements.

Figure 5-7 (Part 1 of 2). Exception Control Statements based on Type 74 (Device Activity) SMF record

Exceptions based on Type 74 (Device Activity) SMF record[9]				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
average response time	DRTAVG	devnmbr, volser, stg grp, or class[10]	SMF74ATV SMF74MEC SMF74SSC SMF74INT SMF74QUE SMF74SAM	(ATV/MEC)+(QUE/SAM)/(SSC/INT)
average device busy delay time	DBDL	devnmbr, volser, stg grp, or class[10]	SMF74DVB SMF74MEC	DVB/MEC
average control unit busy delay time	CUDBL	devnmbr, volser, stg grp, or class[10]	SMF74CUB SMF74MEC	CUB/MEC

[9] specifying a condition based on this record, you must use a qualifier. If a qualifier is not used, a syntax error occurs and the condition is not processed.

[10] devnmbr — specifies a one to three-digit hexadecimal device number in the range 000 through FFF.
volser — specifies a one to six character volume serial number enclosed in quotes.
stg grp — specifies a one to eight character storage group name
class — specifies any of the six valid device classes for Monitor I device activity measurements.

Figure 5-7 (Part 2 of 2). Exception Control Statements based on Type 74 (Device Activity) SMF record

Exceptions based on Type 75 (Page/Swap Data Set Activity) SMF record				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
percent busy-data set	PSBSY	pagename[11]	SMF75USE SMF75SAM	(USE*100)/SAM
average service time per SSCH (swap data sets only)	PSAVG	pagename[11]	SMF75REQ SMF75SAM SMF75SIO SMF75INT	((REQ*INT)/SAM)/SIO
page transfer time (page data sets only)	PSPTT	pagename[11]	SMF75REQ SMF75SAM SMF75PGX SMF75INT	((REQ*INT)/SAM)/PGX
pages transferred per second	PSPT	pagename[11]	SMF75INT SMF75PGX	PGX/INT
page data set activity rate per SECOND	PSART	pagename[11]	SMF75INT SMF75SIO	SIO/INT
average slots used	PSAVGSL	pagename[11]	SMF75AVU	direct comparison
number of bad slots	PSBADS	pagename[11]	SMF75BDS	direct comparison
average swap sets used	PSAVGSE	pagename[11]	SMF75AVU	direct comparison
bad swap sets	PSBADSE	pagename[11]	SMF75BDS	direct comparison

[11] pagename specifies the page or swap data set name, which must conform to data set name standards. Specifying the data set name is optional. Therefore, if the qualifier is omitted, the threshold value applies to each page or swap data set.

Figure 5-8. Exception Control Statements based on Type 75 (Page/Swap Data Set Activity) SMF record

Exceptions based on Type 77 (ENQUEUE Activity) SMF record				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
total enqueue contention time in seconds	ENQT	major/min[12]	SMF77WTT	direct comparison
number of enqueue contention events	ENQNE	major/min[12]	SMF77EVT	direct comparison
average contention time	ENQAVG	major/min[12]	SMF77WTT SMF77EVT	WTT/EVT
maximum contention time	ENQMAX	major/min[12]	SMF77WTX	direct comparison

[12]major/min specifies the one to eight character major name of a resource, optionally followed by a comma and a one to forty-four character minor name. If the qualifier is omitted, the threshold values is checked for every minor name within each major name. If only the major name is specified, the threshold is checked for every minor name within the specified major name. A minor name can not be specified without major name.

Figure 5-9. Exception Control Statements based on Type 77 (Enqueue Activity) SMF record

processors)

Exceptions based on Type 78 (Sub-type 1 - I/O Queuing Activity) SMF record (for 308x, 908x, and 4381 processors):				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
logical control unit activity rate	IOART	lcuid[13]	R781SIS R781INT	SIS/INT
average logical control unit queue length	IOQLEN	lcuid[13]	R781QUE R781TIS	QUE/TIS
percentage of requests deferred	IOD	lcuid[13]	R781TIS R781SIS	100*((TIS-SIS)/TIS)
percentage of requests deferred	IOCUB	lcuid[13]	R781CUB	100*(CUB/TIS)
(CU busy)			R781TIS	
percentage of requests deferred	IODB	lcuid[13]	R781DVB	100*(DVB/TIS)
(device busy)			R781TIS	
percentage all channel paths busy	IOCHPB	lcuid[13]	R781ABY	ABY/SAM

[13]lcuid specifies a three-digit hexadecimal number that identifies a logical control unit. Specifying the qualifier is optional; if it is omitted, the threshold applies to all logical control units in the SMF record.

Figure 5-10. Exception Control Statements based on Type 78 (Sub-type 1 - I/O Queuing Activity) SMF record (for 308x, 908x, and 4381

Exceptions based on Type 78 (Sub-type 3 - I/O Queuing Activity) SMF record (for 3090 processor):				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
I/O processor (IOP) queue activity rate	IOPAC	iopid[14]	R783IQCT SMF78INT	IQCT/INT
I/O processor (IOP) initiative queue average queue length	IOPQL	iopid[14]	R783IQCT R783IQSM	IQSM/IQCT
contention rate of an LCU	IOCTR	lcuid[13]	R783QCT SMF78INT	QCT/INT
average queue length of delayed I/O requests	IODLQ	lcuid[13]	R783QCT R783QSM	QSM/QCT
rate of LCU channel path taken	IOART	lcuid[13]	R783PT SMF78INT	(PTi)/INT[15]
percentage of requests caused by control unit busy	IOCUB	lcuid[13]	R783CUB R783PT	MAX[16](CUBi*100)/PTi[15] MIN[16](CUBi*100)/PTi[15]
percent all channel path busy	IOCHPB	lcuid[13]	R783PB R783GSAM	CHPID(0)*CHPID(1)*CHPID(2)* BUSY BUSY BUSY CHPID(3)[17] BUSY

[13]lcuid specifies a three-digit hexadecimal number that identifies a logical control unit. Specifying the qualifier is optional; if it is omitted, the threshold applies to all logical control units in the SMF record.

[14]iopid specifies a two-digit hexadecimal number that identifies the I/O processor (IOP). Specifying the qualifier is optional.

[15]PTi specifies the sum of each channel path taken where i represents channel path 0 to channel path 3.

[16]MAX/MIN MAX applies to exception operator GE; MIN applies to exception operator LE.

[17]CHPID(i) is calculated as R783PB(i)/R783GSAM.

Figure 5-11. Exception Control Statements based on Type 78 (Sub-type 3 - I/O Queuing Activity) SMF record (for 3090 processor)

Exceptions based on Type 78 (Sub-type 2 - Virtual Storage Activity) SMF record				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm Used To Determine Exception**
maximum amount of SQA expansion into CSA in K bytes	SQAE	none	R782SQEX	direct comparison
minimum size of largest free block - CSA in K bytes	CSAFB	none	R782CSLF	direct comparison
minimum size of largest free block - SQA in K bytes	SQAFB	none	R782SQLF	direct comparison

Figure 5-12. Exception Control Statements based on Type 78 (Sub-type 2 - Virtual Storage Activity) SMF record

EXITS Control Statement

The EXITS control statement specifies whether or not a user exit routine is to be executed during post processing. The syntax of the statement is:

```
EXITS/NOEXITS
```

When EXITS is specified, a user exit routine is executed. See "Report Writing by the Post Processor" in Chapter 7 for information on how to code a user exit routine for the post processor. NOEXITS indicates that no user exit routine is to be executed; NOEXITS is the default.

EXRPTS Control Statement

The EXRPTS control statement is required when you want the post processor to generate interval reports when a particular exception occurs. Each report specified must be separated from any other reports by a comma or a blank. No continuation statements are permitted; however, you can use as many EXRPTS statements as you require. The syntax of the statement is:

```
EXRPTS (control-statement-name(report,report,..))
```

control-statement-name
Specifies a one to eight character name starting with an alphabetic character that associates one or more EXCEPT statements with the EXRPTS statement. The EXCEPT statement defines the exception to RMF; the EXRPTS statement defines the action that RMF is to take.

report
Specifies any one of the Monitor I session reports that are acceptable on the REPORTS control statement. They are listed below:

ALL
CHAN/NOCHAN
CPU/NOCPU
DEVICE(option,option ..)/NODEVICE
ENQ/NOENQ
IOQ/NOIOQ
PAGESP/NOPAGESP
PAGING/NOPAGING
TRACE/NOTRACE
VSTOR(option,list)/NOVSTOR
WKLD(list)/NOWKLD

ALL
Specifies that all Monitor I session reports are to be generated. ALL can be combined with explicit specifications of other options. For example, specifying:

```
EXRPTS(CPUBSY1(ALL))
```

causes all of the Monitor I session reports to be generated.

CHAN/NOCHAN
Specifies whether the channel path activity report is to be generated.

CPU/NOCPU
Specifies whether the processor activity report is to be generated.

DEVICE(option,option..)/NODEVICE
Specifies whether I/O device activity reports are to be generated. You can request device activity by specifying all devices within one or more classes, and, optionally, one or more specific devices.

Any of the following options can be specified:

- a device number in the form NMBR(nmbr1,nmbr.) where nmbr1 and nmbr2 are 3 digit hexadecimal numbers.
- any of the following classes:

<u>CHRDR</u>/NOCHRDR	--	Character reader devices
<u>COMM</u>/NOCOMM	--	Communications equipment
<u>DASD</u>/NODASD	--	Direct access storage devices
<u>GRAPH</u>/NOGRAPH	--	Graphics devices
<u>TAPE</u>/NOTAPE	--	Magnetic tape devices
<u>UNITR</u>/NOUNITR	--	Unit record devices

 LC28-1556-2

- storage groups in the form SG (aaa,bbb) where aaa and bbb are 1 to 8 character names.

When you omit DEVICE and specify ALL, the device classes underscored above are included in the report. When you specify DEVICE, you must include a list of either device classes or addresses or both. When you specify a device class in the option field, the reports generated depend on whether you have used the negative value or the positive value of the option. If you use a negative option, you will get all the device reports with the exception of the option or options you specify. For example, DEVICE(NOTAPE) causes the post processor to generate all I/O device activity reports except the report on magnetic tape devices. If you use a positive option, you will get only the device report corresponding to that option. For example, DEVICE(TAPE) causes the post processor to generate the device activity report for magnetic tape devices; no other I/O device activity reports are printed.

The NMBR field indicates that RMF is to report on the specific devices identified. The numbers can be expressed as a single device or as a range of devices. A range is indicated by specifying the first and last device numbers separated by a colon. Each single number or range is separated by a comma. For example, to request device reporting for magnetic tape devices 180, 183, 184, 185, and 188 as well as all communication equipment, you would specify:

```
DEVICE(COMM,NMBR(180,183:185,188))
```

NONMBR is the default.

The STORAGE GROUP field indicates that RMF is to report on the specific storage groups identified. The storage groups can be expressed as a single storage group name or as a range of storage groups. A range is indicated by specifying the first and last storage group names separated by a colon. Each single storage group name or range is separated by a comma. NOSG is the default and it causes any existing list of SG names to be cancelled.

ENQ/NOENQ
Specifies whether the enqueue activity interval report is to be generated. The level of enqueue activity reporting depends on the level selected during the reporting period.

IOQ/NOIOQ
Specifies whether the I/O queuing activity report is to be generated.

PAGESP/NOPAGESP
Specifies whether the page/swap data set activity report is to be generated.

PAGING/NOPAGING
Specifies whether the session paging activity report is to be generated.

TRACE/NOTRACE
Specifies whether the trace activity interval report is to be generated.

<u>VSTOR</u>[({<u>S</u> | D} [,jobname1,jobname2,..])] /NOVSTOR
Specifies whether the virtual storage activity report is to be generated. RMF can produce common storage summary and detail reports and private area summary and detail reports. When you specify S, either explicitly or by default, RMF produces summary reports; when you specify D, RMF produces both summary reports and detail reports.

The Monitor I session gathers private area data only when you specify a jobname on the VSTOR option during the session. The post processor, however, reports any private area data that it finds in the input records. Thus, it is not necessary to identify specific jobnames for the post processor. (If you identify a specific jobname, the post processor produces a private area report for that job only, and only if private area data for it exists in the input records.) It is, indeed, a good practice to omit specific jobnames on the post processor control statements. This practice enables you to use the same post processor control statement to obtain common storage report(s) or to obtain both common storage report(s) and private area report(s) when data exists for private area report(s).

If you specify VSTOR without any operands, RMF produces a summary report for common storage. Examples of other possible combinations are:

- EXRPTS(VSTOR(D)) produces a summary and detail report for common storage. The post processor also produces a summary and detail report for any private area data in the input records.
- EXRPTS(VSTOR(D,VTAM)) produces a summary and detail report for common storage and a summary and detail report for the private area of the VTAM address space. The post processor does not produce reports for any other private area data in the input records.
- EXRPTS(VSTOR(MYJOB)) produces a summary report for common storage and a summary report for the private area of the MYJOB address space. The post processor does not produce reports for any other private area data in the input records.

WKLD(list)/NOWKLD

Specifies whether the system workload activity report is to be generated and indicates the types of sub-reports to be included.

When WKLD is specified, list must indicate the types of sub-reports to be included. Any or all of the following can be specified in the list:

PERIOD/NOPERIOD	-- Requests reporting by performance group period (type 1 sub-report)
GROUP/NOGROUP	-- Requests reporting by performance group (type 2 sub-report)
RANGE/NORANGE	-- Requests reporting by performance objective, domain, and performance group in a range of performance group numbers (type 3, 4, and 5 sub-reports)
OBJ/NOOBJ	-- Requests reporting by performance objective (type 6 sub-report)
DOMAIN/NODOMAIN	-- Requests reporting by domain number (type 7 sub-report)
SYSTEM/NOSYSTEM	-- Requests reporting by system (type 8 sub-report)
TIME/NOTIME	-- Requests reporting by time slice group (type 9 sub-report)

When you omit WKLD and specify ALL, the underscored reports are generated. When you specify WKLD, you must include **list** When you specify a report type in the **list** field, the reports generated depend on whether you have used the negative value or the positive value of the option. If you use the negative value of the option, you get all the workload reports except the options you specify. For example, WKLD(NODOMAIN) causes the post processor to generate all the workload activity reports except the reports by domain number. If you use the positive value of the option, you get only the workload report for that option. For example, WKLD(DOMAIN) causes the post processor to generate only the domain sub-report; no other workload activity reports are printed.

Example: The post processor is to generate a channel path activity report for each interval when the I/O service rate for performance group number 3 is less than or equal to 100 service units per second. Use the following control statements:

```
EXCEPT(CHNRPT(IOSRV(0030),LE,100))
EXRPTS(CHNRPT(CHAN)))
```

MAXPLEN Control Statement

The MAXPLEN control statement specifies the maximum number of lines of plotted data that can appear in a single plot report. The syntax of the statement is:

```
MAXPLEN(nnn)
```

where **nnn** is the number of lines. The maximum value allowed is 999. When you omit the MAXPLEN statement, the default number of lines is 50. When you explicitly specify MAXPLEN, you must specify the number of lines.

When a plot report exceeds one page, the plot lines run continuously from one page to the next to preserve the unity of the plotted data; there is no break or repetition of headings on each successive page.

PINTV Control Statement

The PINTV control statement specifies the length of time that is to elapse between each line of plotted data, that is, the length of the plot interval. The syntax of the statement is:

```
PINTV(hhmm)
```

Where **hh** is the hour and **mm** is the minute on a 24-hour clock. Note that the time is the length of the plot interval, not the time of day when the plot reporting period begins.

When you omit the PINTV statement but indicate that plot reports are required by specifying the PLOTS control statement, the post processor generates a line in a plot report for each RMF interval included in the reporting period. Thus, the default for the PINTV control statement is the length of the RMF measurement interval in effect during the reporting period. If you explicitly specify PINTV, you must explicitly specify a time value. Specifying PINTV without a time value causes a syntax error. Note that the PINTV statement is ignored when PLOTS is not specified.

In deciding the length of the plot interval, consider its relationship to the length of the RMF measurement interval. When you specify a length that is less than that of the measurement interval, the post processor substitutes the length of the measurement interval. When you specify a length that is greater than that of the measurement interval, the post processor combines data collected during all the measurement intervals that fall within the plot interval.

Thus, specifying a length less than the measurement interval has no effect. For most effective plot reporting, the plot interval should be either equal to or a multiple of the length of the RMF measurement interval. Specifying a plot interval that is a multiple of the measurement interval is useful when you want a very high-level plot of system activity.

Note: Because the length of the RMF measurement intervals can vary by a few seconds or more, the number of record intervals that the post processor includes in a plot interval can also vary.

PLOTS Control Statement

The PLOTS control statement specifies that plot reports are to be produced and identifies the plots that you require. The syntax of the statement is:

```
PLOTS(plot[,plot],..)
```

where **plot** can be any of the following:

BATCH	--	Maximum number of batch users produced for all channels on all processors.
CH(nn)	--	Channel path busy percentage, where **nn** is the hexadecimal channel path identifier. When CH is requested, the channel path identifier must be specified.
CPUID[(n)]	--	Processor busy percentage, where **n** is the processor identifier. When **n** is omitted, plots are produced for all processors.
CSAA	--	Maximum allocated CSA (below the 16-megabyte line).
CSAFP	--	Minimum number of free CSA pages (below the 16-megabyte line).
CUBDL(nnn)	--	Average number of milliseconds when any I/O requests for the device was delayed because the control unit were busy, where **nnn** is the hexadecimal device number of the device to be included in the report. When the control unit busy delay time plot is requested, **nnn** must be specified to identify the device.
DAC(nnn)	--	Device activity rate, where **nnn** is the hexadecimal device number of the device to be included in the report. When the device activity plot is requested, **nnn** must be specified to identify the device.
DBDL(nn)	--	Average number of milliseconds when any I/O requests for the device were delayed because the device was busy, where **nnn** is the hexadecimal device number of the device to be included in the report. When the device busy delay time plot is requested, **nnn** must be specified to identify the device.
DCON(nnn)	--	Device connect time, where **nnn** is the hexadecimal device number. When device connect time is requested, **nnn** must be specified to identify the device.
DRT(nnn)	--	Device response time where **nnn** is the hexadecimal device number. When device response time is requested, **nnn** must be specified to identify the device.
ESMR	--	Average rate of page migration from extended to auxiliary storage.
HUIC(nn)	--	Average high unreferenced interval count.
IOAC(nnn)	--	Rate of requests successfully initiated for the logical control unit, where **nnn** is the hexadecimal identifier of the logical control unit. When IOAC is requested, **nnn** must be specified to identify the logical control unit.
PAGE	--	Demand paging rate.
PTES	--	Average paging transfer rate to extended storage.
SEV	--	Service rate.
SQAA	--	Maximum allocated SQA (below the 16-megabyte line).
SQAE	--	Amount of SQA expansion into CSA (below the 16-megabyte line).
SQAFP	--	Minimum number of free SQA pages (below the 16-megabyte line).
STC	--	Maximum number of started task and mount task users.
SWA	--	Swap rate SWA option.
TPAG	--	Total paging rate.
TRA	--	Transaction rate.
TSO	--	Maximum number of TSO sessions.

Each type of plot is described in more detail in "Plot Reports" later in this chapter.

When you omit the PLOTS statement, no plot reports can be produced. If you specify PLOTS, you must identify at least one plot type. No continuation statements are permitted for the PLOTS statement. When you require more plot types than can fit on one statement, specify the additional plot types by supplying additional PLOTS statements. There is no limit on the number of PLOTS statements. For example:

```
PLOTS(DRT(230),DCON(230),CH(04))
PLOTS(DAC(350))
PLOTS(STC,SEV)
PLOTS(PAGE,TPAG,SWA,CPUID)
```

Note: The order in which you specify the plot types does not affect the order in which RMF generates the requested plot reports.

PTOD Control Statement

The PTOD control statement specifies the starting and ending time of the reporting period for a plot report for each day in the reporting period. The syntax of the statement is:

```
PTOD(hhmm,hhmm)
```

where **hh** is the hour and **mm** is the minute on a 24-hour clock. The first time specifies the beginning of the reporting period and the second time specifies the end of the reporting period. The second time you specify must be later than the first, or a syntax error occurs.

When the PTOD statement is omitted, the default value is PTOD(0000,2400); that is, all times are reported. Thus, you would use this control statement when you want a reporting period for plot reports that is different from the default value.

Because the range of values allowed is from 0000 to 2400, it is not possible to define a reporting period that consists of a single block of time that spans more than one calendar day. Thus, you cannot define a reporting period that runs from 12 noon on one day to 12 noon on the next day. However, you can define a reporting period that consists of the same block of time over several days. For example, to produce a plot report using data collected from 8:00 A.M. to 4:00 P.M. for the week beginning on January 3, 1982, and ending on January 9, 1982, the required DATE and PTOD statements would be:

```
DATE(82003,82009)
PTOD(0800,1600)
```

REPORTS Control Statement

The REPORTS control statement specifies the reports to be generated by the post processor. This statement is required to generate duration reports or interval reports. If you specify DINTV, duration reports are generated. If you do not specify DINTV, interval reports are generated. Each report specified must be separated from any other reports by a comma or a blank. No continuation statements are permitted; however, you can use as many REPORTS statements as you require. The syntax of the statement is:

```
REPORTS(report[,report],..)
```

where **report** can be any of the options described later in this section. ALL indicates that all reports are to be generated; the option is described in detail later in this section.

When you specify an option either alone or in combination with ALL, the reports generated depend on whether you have used the negative value or the positive value of the option. This concept can perhaps best be illustrated by the following examples:

- When you specify REPORTS(CPU), the only report generated is the processor activity report.
- When you specify either REPORTS(NOCPU) or REPORTS(ALL,NOCPU), the post processor generates all reports except for processor activity.
- When you specify REPORTS(ALL,CPU) or REPORTS(ALL), the post processor generates all reports.

Notes:

1. There is no additional overhead introduced when you specify REPORTS(ALL) when some system activities were not measured during the reporting period.
2. When the DINTV statement is present, all options except ALL, CPU, CHAN, DEVICE, WKLD, PAGESP, IOQ, PAGING, or VSTOR are ignored. These options are described in the following list.
3. Regardless of the specification on the REPORTS statement, the post processor can generate a report on a system activity only when that activity was measured during the reporting period and the RECORD option was specified.

ALL

Specifies that all Monitor I and Monitor II reports are to be generated. Any user-supplied Monitor II background session reports are also included when ALL is specified. ALL can be combined with explicit specifications of other options. For example, specifying:

```
REPORTS(ALL,NOENQ,DEVICE(NOUNITR,NOCOMM))
```

would cause all of the reports to be generated, with the exception of enqueue activity and device activity for unit record devices and communication equipment.

When ALL is specified on the REPORTS statement and the DINTV statement is present, duration reports are produced for processor, channel path, device, workload, page/swap data set, I/O queuing, and paging activity.

ARD[(class,status,domain)]/NOARD

Specifies whether the Monitor II address space resource data report is to be generated, where the **class, status,** and **domain** operand fields specify the selection criteria for the address spaces to be included. The operand fields are positional; thus, any operand omitted, except the last, must be replaced by a comma. When you do not specify a value for **class, status,** or **domain**, the menu default, if present, is used; if there is no menu default, the value used to generate the SMF record is used to generate the printed report. The meaning of each operand field is:

class

A	specifies all address spaces
B	specifies batch, started task, and mount task address spaces
T	specifies TSO address spaces

status

A	specifies all address spaces
I	specifies address spaces that are active

domain

A	specifies all address spaces
nnn	specifies those address spaces in the domain identified by **nnn**

ARDJ[(jobname)]/NOARDJ
Specifies whether the Monitor II address space resource data by jobname report is to be generated, where **jobname** identifies a specific job for the report. When you omit **jobname**, an address resource data by jobname report is generated for each job for which address space resource data by jobname was collected.

ASD[(class,status,domain)]/NOASD
Specifies whether the Monitor II address space state data report is to be generated. The meaning and use of **class, status**, and **domain** are as described above under ARD.

ASDJ[(jobname)]/NOASDJ
Specifies whether the Monitor II address space state data by jobname report is to be generated, where **jobname** identifies a specific job for the report. When you omit **jobname**, an address space state data by jobname report is generated for each job for which address space state data by jobname was collected.

ASRM[(class,status,domain)]/NOASRM
Specifies whether the Monitor II address space SRM data report is to be generated. The meaning and use of **class, status,** and **domain** are as described above under ARD.

ASRMJ[(jobname)]/NOASRMJ
Specifies whether the Monitor II address space SRM data by jobname report is to be generated, where **jobname** identifies a specific job for the report. When you omit **jobname**, an address space SRM data by jobname report is generated for each job for which address space SRM data by jobname was collected.

CHAN/NOCHAN
Specifies whether the Monitor I channel path activity report is to be generated.

CHANNEL /NOCHANNEL
Specifies whether the Monitor II channel path activity report is to be generated.

CPU/NOCPU
Specifies whether the Monitor I processor activity report is to be generated.

DDMN/NODDMN
Specifies whether the Monitor II domain activity report is to be generated.

DEV [(type)]/NODEV
Specifies whether the Monitor II table report for I/O device activity is to be generated. You can request device activity by specifying all devices in one class, or one or more specific device numbers, volume serial numbers, or storage groups.

One of the following types can be specified:

- A device class:

DASD	Direct access storage devices
DASD	Direct access storage devices
TAPE	Magnetic tape devices
COMM	Communication equipment
CHRDR	Character reader devices
UNITR	Unit record devices
GRAPH	Graphic devices

- One or more volume serial numbers:

```
          / { aaaaaa               } \
{ VOLSER } ( { aaaaaa,bbbbbb:zzzzzz } )
{ V      }  \ { aaaaaa,bbbbbb,..     } /
```

- One or more device numbers:

```
          / { aaa         } \
{ NUMBER } ( { aaa,bbb:zzz } )
{ N      }  \ { aaa,bbb,..  } /
```

- One or more storage group names:

```
        / {  aaaaaaaa                   } \
{ SG } ( {  aaaaaaaa,bbbbbbbb,...       } )
{ S  }  \ {  aaaaaaaa,bbbbbbbb:zzzzzzzz } /
```

DEVICE(option,option..)/NODEVICE

Specifies whether Monitor I session I/O device activity reports are generated. You can request device activity by specifying all devices within one or more classes, and, optionally, one or more specific devices.

Any of the following options can be specified:

- a device number in the form NMBR(nmbr1,nmbr.) where nmbr1 and nmbr2 are 3 digit hexadecimal numbers.
- any of the following classes:

<u>CHRDR</u>/NOCHRDR	-- Character reader devices
<u>COMM</u>/NOCOMM	-- Communications equipment
<u>DASD</u>/NODASD	-- Direct access storage devices
<u>GRAPH</u>/NOGRAPH	-- Graphics devices
<u>TAPE</u>/NOTAPE	-- Magnetic tape devices
<u>UNITR</u>/NOUNITR	-- Unit record devices

- storage groups in the form SG (aaa,bbb) where aaa and bbb are 1 to 8 character names.

Note: Because of its speed, the IBM 3890 Document Processor is reported under unit record devices rather than character reader devices.

When you omit DEVICE and specify ALL, the device classes underscored above are included in the report. When you specify DEVICE, you must include a list of either device classes, numbers, or both. When you specify a device class in the **option** field, the reports generated depend on whether you have used the negative value or the positive value of the option. If you use a negative option, you will get all the device reports with the exception of the option or options you specify. For example, DEVICE(NOTAPE) causes the post processor to generate all I/O device activity reports except the report on magnetic tape devices. If you use a positive option, you will get only the device report corresponding to that option. For example, DEVICE(TAPE) causes the post processor to generate the device activity report for magnetic tape devices; no other I/O device activity reports are printed.

The NMBR field indicates that RMF is to report on the specific devices identified. The numbers can be expressed as a single device or as a range of devices. A range is indicated by specifying the first and last device numbers separated by a colon. Each single number or range is separated by a comma. For example, to request device reporting for magnetic tape devices 180, 183, 184, 185, and 188 as well as all direct access devices and communication equipment, you would specify:

```
DEVICE(COMM,DASD,NMBR(180,183:185,188))
```

RMF reports on the storage groups you specify in the SG field. You can select one storage group name or a range of storage groups. To select a range of storage groups, specify NODASD, and the first name and the last name with a semicolon between them. For example to select the range of storage groups from PROC01 to PROC05, specify:

```
DEVICE(NODASD,SG(PROC01:PROC05))
```

To select one storage group, for example, PROC02, specify:

```
DEVICE(NODASD,SG,(PROC02))
```

DEVV [(id)]/NODEVV

Specifies whether the Monitor II row report for I/O device activity is to be generated. You can request device activity for a specific device by specifying either:

{VOLSER | V} (xxxxxx)

Requests I/O device activity for the specific volume identified by the volume serial number xxxxxx.

{NUMBER | N} (yyy)

Requests I/O device activity for the specific device identified by the number yyy.

RMF reports a storage group name for each volume assigned to one.

ENQ/NOENQ

Specifies whether the Monitor I enqueue activity interval report is to be generated. The level of enqueue activity reporting depends on the level selected at the time the data was gathered.

IOQ/NOIOQ

Specifies whether a Monitor I session I/O queuing activity report is to be generated.

IOQUEUE/NOIOQUEUE

Specifies whether a Monitor II I/O queuing activity report is to be generated.

PAGESP/NOPAGESP

Specifies whether the Monitor I page/swap data set activity report is to be generated.

PAGING/NOPAGING

Specifies whether the Monitor I paging activity report is to be generated.

PGSP [**(PAGE)** | **(SWAP)**] **/NOPGSP**

Specifies whether the Monitor II page/swap data set activity report is to be generated.

SENQ [**(S)** | **(D)** | **(A,sysname)** | **(E,sysname)** | **(majorname[,minorname])**] **/NOSENQ**

Specifies whether the Monitor II enqueue activity report is to be generated. The operands describe the type of data you require. Only one operand can be specified. The meaning of each operand field is:

S

Specifies a summary report. For each resource that had contention activity, the summary report includes the scope of the resource, the number of tasks that own the resource, the number of tasks waiting for exclusive use of the resource, and the number of tasks waiting for shared use of the resource. If you specify SENQ without operands, S is the default.

Note: If the session option specified a resource or group of resources by name, the report includes data for only those resources.

D

Specifies a detail report. For each resource that had enqueue contention, the detail report includes the scope of the resource, the name and address space identifier of each job owning or waiting for the resource, and the type and status of each job's request for the resource.

Note: If the session option specified a resource or group of resources by name, the report includes data for only those resources.

A,sysname

Specifies a report that includes all resources that a specific system holds in a global resource serialization complex, where **sysname** indicates the system for which the report is requested.

E,sysname

Specifies a report that includes all exclusively - owned resources that a specific system held in a global resource serialization complex, where **sysname** indicates the system for which the report is requested.

majorname[,minorname]

Specifies a detail report for a specific resource. The **majorname** field, which corresponds to the **qname** field in the ENQ and DEQ macro instructions, contains the one to eight character major name of a serially-reusable resource. Optionally, the major name is followed by a comma and a minor name. The **minorname** field, which corresponds to the **rname** field in the ENQ and DEQ macro instructions, contains the minor name of the resource.

The maximum length of the field is 32 characters, including the comma. Because the major name is 1 to 8 characters, the minor name can be from 1 to 30 characters, depending on the length of the major name. If you want a report on a minor name, but the **majorname,minorname** operand exceeds 32 characters, you must specify only the major name. RMF then reports on all resources grouped under the major name.

Notes:

1. RMF treats the single character A, D, E or S as a request for the report. Therefore, A or E cannot be used as a major name; S or D cannot be used as a major name unless a minor name is also specified.

2. If the session option identified a different resource or group of resources, RMF issues a message to tell you that no data was available to meet your selection criteria.

SENQR [**(ALLVSER)** | **(volser)**] **/NOSENQR**

Specifies whether the Monitor II reserve activity report is to be generated. The operands describe the type of data you require. Only one operand can be specified. The meaning of each operand field is:

ALLVSER

Specifies a report that describes all reserve requests. If you specify SENQR without operands, ALLVSER is the default.

Note: If the session option identified a specific device by volume serial number, the report includes data for only that device.

volser

Specifies a report that describes the reserve requests for a particular device, where **volser** is the one to six character volume serial number of the volume mounted on the device.

Note: If the session option identified a different specific device, RMF issues a message to tell you that no data was available to meet your selection criteria.

SRCS/NOSRCS

Specifies whether the Monitor II real storage/processor/SRM report is to be generated.

SPAG/NOSPAG

Specifies whether the Monitor II session paging activity report is to be generated.

TRACE/NOTRACE

Specifies whether the Monitor II trace activity interval report is to be generated.

TRX [(sname,nnn[,nnn,nnn:nnn],ALLPGN)]/NOTRX

Requests the Monitor II transaction activity report. Because the report requires data that the SRM collects only when RMF is monitoring workload activity, the report can be obtained only when a Monitor I session was active with the WKLD option in effect during the reporting period. Otherwise, RMF issues a message to inform you that the report is unavailable because a Monitor I session with the WKLD option was not active.

The meaning of each operand field is:

sname

Requests data for all performance groups associated with the subsystem identified by **sname**, where **sname** is a 1-4 character subsystem name that is defined in the installation control specification.

nnn[,nnn,nnn:nnn]

Requests data for one or more performance groups, where **nnn** is a 1-3 digit performance group number. Performance groups can be specified singly (nnn) or as a range (nnn:nnn), where the lower and upper bounds are separated by a colon. Each number or range of numbers must be separated by a comma.

ALLPGN

Requests data for all performance groups. ALLPGN is the default.

The operands are not positional; thus, an omitted operand does not require a comma in its place. More than one operand can be specified. While there is no reason to combine ALLPGN with any other operand, you can specify on the same request both one or more subsystem names and individual performance groups or ranges of performance groups. In this case, if you specify individually a performance group that is also associated with a subsystem, data for that performance group appears twice in the report.

The menu default for the maximum number of performance groups that can be reported in response to any one request is 504.

user-report[(operands)]

Specifies that a user-supplied Monitor II session report is to be generated, where **user-report** is the name of the option used to collect data for the report and **operands** are any operands your installation established when the report was designed. When your report has operands but you do not specify any operands on the REPORTS control statement, the post processor uses the menu default, if present, for any omitted operand. When there is no menu default, the post processor takes the operand in effect when the data was collected.

See "Monitor II Session User Reports" in Chapter 7 for a description of how to add a user-supplied report to the post processor. After you have performed the steps that make your report available to the post processor, your report will be printed when you specify the name of the report or ALL on the REPORTS control statement.

VSTOR[({S | D}[,jobname1,jobname2,..])] /NOVSTOR

Specifies whether the virtual storage activity report is to be generated. RMF can produce common storage summary and detail reports and private area summary and detail reports. When you specify S, either explicitly or by default, RMF produces summary reports; when you specify D, RMF produces both summary reports and detail reports.

The Monitor I session gathers private area data only when you specify a jobname on the VSTOR option during the session. The post processor, however, reports any private area data that it finds in the input records. Thus, it is not necessary to identify specific jobnames for the post processor. (If you identify a specific jobname, the post processor produces a private area report for that job only, and only if private area data for it exists in the input records.) It is, indeed, a good practice to omit specific jobnames on the post processor control statements. This practice enables you to use the same post processor control statement to obtain common storage report(s) or to obtain both common storage report(s) and private area report(s) when data exists for private area report(s).

If you specify VSTOR without any operands, RMF produces a summary report for common storage. Examples of other possible combinations are:

- REPORTS(VSTOR(D)) produces a summary and detail report for common storage. The post processor also produces a summary and detail report for any private area data in the input records.
- REPORTS(VSTOR(D,VTAM)) produces a summary and detail report for common storage and a summary and detail report for the private area of the VTAM address space. The post processor does not produce reports for any other private area data in the input records.
- REPORTS(VSTOR(MYJOB)) produces a summary report for common storage and a summary report for the private area of the MYJOB address space. The post processor does not produce reports for any other private area data in the input records.

WKLD(list)/NOWKLD

Specifies whether the Monitor I system workload activity report is to be generated and indicates the types of sub-reports to be included. When the DINTV statement is present, a duration report is produced; otherwise, an interval report is produced.

When WKLD is specified, **list** must indicate the types of sub-reports to be included. Any or all of the following can be specified in the list:

PERIOD/NOPERIOD	--	Requests reporting by performance group period (type 1 sub-report)
GROUP/NOGROUP	--	Requests reporting by performance group (type 2 sub-report)
RANGE/NORANGE	--	Requests reporting by performance objective, domain, and performance group in a range of performance group numbers (type 3, 4, and 5 sub-reports)
OBJ/NOOBJ	--	Requests reporting by performance objective (type 6 sub-report)
DOMAIN/NODOMAIN	--	Requests reporting by domain number (type 7 sub-report)
SYSTEM/NOSYSTEM	--	Requests reporting by system (type 8 sub-report)
TIME/NOTIME	--	Requests reporting by time slice group (type 9 sub-report)

When you omit WKLD and specify ALL, the underscored reports are generated. When you specify WKLD, you must include **list** When you specify a report type in the **list** field, the reports generated depend on whether you have used the negative value or the positive value of the option. If you use the negative value of the option, you get all the workload reports except the options you specify. For example, WKLD(NODOMAIN) causes the post processor to generate all the workload activity reports except the reports by domain number. If you use the positive value of the option, you get only the workload report for that option. For example, WKLD(DOMAIN) causes the post processor to generate only the domain sub-report; no other workload activity reports are printed.

RTOD Control Statement

The RTOD control statement specifies the starting time and ending time of the reporting period for interval or duration reporting for each day included in the reporting period. The syntax of the statement is:

```
RTOD(hhmm,hhmm)
```

where **hh** is the hour and **mm** is the minute on a 24-hour clock. The first time specifies the beginning of the reporting period and the second time specifies the end of the reporting period. The second time must be later than the first, or a syntax error occurs. When the RTOD statement is omitted, the default value is RTOD (0000,2400); that is, all times are reported. Thus, you would use this statement when you want a reporting period for interval or duration reports that is different from the default value.

Because the range of values allowed is from 0000 to 2400, it is not possible to define a reporting period that consists of a single block of time that spans more that on calendar day. Thus, for example, you cannot define a reporting period that runs from 12 noon on one day to 12 noon on the next day. However, you can define a reporting period that consists of the same block of time over several days. For example, to produce interval reports including data collected every morning from 8:00 A.M. to 1:00 P.M. for the week beginning on January 3, 1982, and ending on January 9, 1982, the required DATE and RTOD statements would be:

```
DATE(82003,82009)
RTOD(0800,1300)
```

SESSION Control Statement

The SESSION control statement specifies the particular Monitor II background session that created the SMF records to be included in the reports. The syntax of the statement is:

```
SESSION(session-id)
```

where **session-id** is the two-character alphameric session identifier of the particular session. If you explicitly specify SESSION, you must supply a session identifier. Only one session-id may be reported on during a post processing session.

When you omit the SESSION statement, all SMF records that fall within the reporting period and are pertinent to the types of reports specified on the REPORTS statement are included in the reports, regardless of the session that created them.

STOD Control Statement

The STOD control statement specifies the starting time and ending time of the reporting period for a summary report for each day in the reporting period.

The syntax of the statement is:

```
STOD(hhmm,hhmm)
```

where **hh** is the hour and **mm** is the minute on a 24-hour clock. The first time specifies the beginning of the reporting period and the second time specifies the end of the reporting period. The second time must be later than the first, or a syntax error occurs.

When the STOD statement is omitted, the default value is STOD(0000,2400); that is, all times are reported. Thus, you would use this control statement when you want a reporting period for a summary report that is different from the default value.

Because the range of values allowed is from 0000 to 2400, it is not possible to define a reporting period that consists of a single block of time that spans more than one calendar day. Thus, you cannot define a reporting period that, for example, runs from 12 noon on one day to 12 noon on the next day. However, you can define a reporting period that consists of the same block of time over several days. For example, to produce a summary report using data collected from 8:00 A.M. to 1:00 P.M. for the week beginning on January 3, 1982, and ending on January 9, 1982, the required DATE and STOD statements would be:

```
DATE(82003,82009)
STOD(0800,1300)
```

SUMMARY Control Statement

The SUMMARY control statement specifies whether a summary report is to be produced and indicates the type of summary data that you require. The syntax of the statement is:

SUMMARY(type)/NOSUMMARY

where **type** can be either or both of the following:

INT -- indicating that one interval summary line is to be produced for each measurement interval that falls within the reporting period.

TOT -- indicating that one total summary data line is to be produced for all the measurement intervals that fall within the reporting period.

When both are specified, INT and TOT can appear in any order. When you explicitly specify SUMMARY, you must specify the type of summary data that you require. Specifying SUMMARY without **type** causes a syntax error. When you omit the SUMMARY statement, the default is SUMMARY(INT,TOT). That is, a summary report is produced, and the report includes both interval summary data lines and a total summary data line.

When a summary report consists of more than one page, the headings are repeated for each page. When total summary data is requested, a total summary line is generated for the intervals on each page, and the last page of the report contains a total summary data line that reflects the contents of all pages in the report.

SYSID Control Statement

The SYSID control statement specifies the four-character system identifier of the system about which reports should be generated. The syntax of the statement is:

```
SYSID(cccc)
```

where **cccc** can be any four alphameric and/or special characters and specifies the system identifier. When you explicitly specify SYSID, you must supply the system identifier. You can only specify one SYSID control statement per post processing session. Specifying the SYSID control statement causes the post processor to include in the reporting all pertinent SMF records that have a matching system identifier.

Omitting SYSID causes the post processor to include in the reporting all SMF records, for all system identifiers. When more than one system identifier is encountered, the post processor produces separate reports for each system encountered.

SYSOUT Control Statement

The SYSOUT control statement specifies the SYSOUT class for all formatted report output. The syntax of the statement is:

```
SYSOUT(class)
```

where **class** is the desired SYSOUT class. When you explicitly specify SYSOUT, you must indicate a SYSOUT class. When you omit the SYSOUT statement, the default is SYSOUT class A. The SYSOUT class for post processor messages is always A; it is not affected by the SYSOUT control statement. The message SYSOUT class can be changed by preallocating MFPMSGDS.

Examples of Control Statements

The examples in this section show various uses of the post processor. All the examples include the DATE statement to illustrate how the value specified for DATE relates to the value specified for RTOD, PTOD, ETOD, or STOD to define the reporting period. During actual execution of the post processor, your installation might find it more useful to control the dates included in the reports by controlling the contents of the input data set and omitting the DATE statement. Because the default for the DATE statement is a reporting period that encompasses all dates included in the SMF records in the input data set, omitting the DATE statement enables you to establish a set of control statements that can be used on a regular schedule without modification.

1. The post processor is to generate all interval reports except tracing. The reporting period runs from 8:00 A.M. to 12 noon for the five-day business week from January 4, 1982 to January 8, 1982. Use the following control statements:

   ```
   DATE(82004,82008)
   REPORTS(ALL,NOTRACE)
   RTOD(0800,1200)
   NOSUMMARY
   ```

   Because the EXITS, SYSID, and SYSOUT statements are omitted and their defaults taken, no user exits are entered, all systems are included in the reports, and any report and message output is sent to SYSOUT class A.

2. The post processor is to generate plot reports of the maximum number of batch users, the processor busy percentage, and the device activity rate for an I/O device that has an address of 14E. The reporting period is the twelve-hour period beginning at 12:00 A.M. on January 5, 1982. Use the following control statements:

   ```
   DATE(82005,82005)
   MAXPLEN(100)
   PINTV(0015)
   PLOTS(BATCH,CPUID,DAC(14E))
   PTOD (0000,1200)
   NOSUMMARY
   ```


Notes:

- The MAXPLEN statement limits the number of plot lines to 100.
- The PINTV statement sets the plot interval -- the time to elapse between lines in the plot report -- to 15 minutes. If the RMF interval length in effect during the reporting period was identical to the length desired for the plot interval, the PINTV statement could be omitted. The default for the PINTV statement is the length of the RMF measurement interval.
- Because no processor identifier is specified on the CPUID plot request, a plot is produced for each processor encountered.
- Because the EXITS, SYSID, and SYSOUT statements are omitted and their defaults taken, no user exits are entered, all systems are included in the reports, and any report and message output is sent to SYSOUT class A.

3. The post processor is to generate duration reports for processor activity, channel path activity, and I/O device activity for magnetic tape devices, direct access devices, and communications equipment. The reporting period is the twelve-hour period from 6:00 A.M. to 6:00 P.M. on June 22, 1982. The duration interval is six hours, causing two duration reports to be produced for each specified activity. Use the following control statements:

```
DATE(82173,82173)
DINTV(0600)
REPORTS(CPU,CHAN)
REPORTS(DEVICE(TAPE,DASD,COMM))
RTOD(0600,1800)
NOSUMMARY
```

   Because the EXITS, SYSID, and SYSOUT statements are omitted and their defaults taken, no user exits are entered, all systems are included in the reports, and any reports, and message output is sent to SYSOUT class A.

4. To obtain a summary report, including both interval summary data and total summary data, based on SMF record data collected from 9:00 A.M. to 5:00 P.M. for the five-day business week from June 4, 1982, to June 8, 1982, use the following control statements:

```
DATE(82165,82169)
SUMMARY(INT,TOT)
STOD(0900,1700)
```

   Notes:

   - The SUMMARY(INT,TOT) statement could be omitted because the default for the SUMMARY control statement is a summary report that includes both interval summary data and total summary data.
   - Because the EXITS, SYSID, and SYSOUT statements are omitted and their defaults taken, no user exits are entered, all systems are included in the report, and any report and message output is sent to SYSOUT class A.

5. Assume that a single execution of the post processor is to generate the following types of reports:

   - Interval reports for all Monitor I session activities except tracing. The reporting period is from 8:00 A.M. to 12 noon on June 18, 1982.
   - Interval reports for Monitor II address space state data, address space resource data, and address space SRM data. No other Monitor II session interval reports are required. The data for the reports is the data collected by the Monitor II session with the session identifier of BB. The reporting period is from 8:00 A.M. to 12 noon on June 18, 1982. The reports are to reflect changed values; that is, the reports are to be in delta mode.
   - Plot reports of the maximum number of batch users, the maximum number of started task and mount task users, and the maximum number of TSO sessions. The reporting period is from 8:00 A.M. to 12 noon on June 18, 1982. The plot interval is 30 minutes, and the maximum number of lines per plot is 20.
   - A summary report covering each interval in the reporting period. The reporting period is from 8:00 A.M. to 12:00 noon on June 18, 1982.

   The control statements required to obtain the output reports defined above are:

```
DATE(82169,82169) - establishes date of reporting period
RTOD(0800,1200) - establishes time of reporting period for
                    interval reports
REPORTS(ALL,NOTRACE,NOARDJ,NOASDJ)
REPORTS(NOCHANNEL,NODDMN,NODEV)        defines interval
REPORTS(NODEVV,NOIOQUEUE,NOPGSP)       reports to be
REPORTS(NOSENQ,NOSENQR,NOTRX)          generated
REPORTS(NOASRMJ,NOSRCS,NOSPAG)
SESSION(BB)
               defines session and mode for Monitor II
               session interval reports
DELTA
PLOTS(BATCH,STC,TSO)
PTOD(0800,1200)
                          defines plot reports
                          to be generated
MAXPLEN(20)
PINTV(0030)
STOD(0800,1200)
                          defines summary report
                          to be generated
SUMMARY(INT)
```

   Because the EXITS, SYSID, and SYSOUT statements are omitted and their defaults taken, no user exits are entered, all systems are included in the reports, and any report and message output is sent to SYSOUT class A.

6. The post processor is to generate the following:

   a. A line in the exception report when the I/O service rate for performance group number 3 is less than or equal to 100 service units per second and the percent device utilization for device 6D8 is greater than or equal to 3.

      A channel path report and a device activity report for the DASD device class if both conditions are met.

      The statements indicated by **(1)** describe the conditions listed in this item.

   b. A line in the exception report if the percent busy for channel path 01 is greater than or equal to ten percent.

      The statement indicated by **(2)** describes the condition listed in this item.

   c. A line in the exception report if the percent busy for processor 01 is less than or equal to 80 percent busy and the average number of TSO users is greater than ten.

      The statements indicated by **(3)** describe the conditions listed in this item.

      The reporting period is the eight-hour interval from 8:00 AM to 4:00 PM for the week beginning June 14, 1982 and ending June 20, 1982.

      Use the following control statements:

```
     DATE(82165,82171)
     ETOD(0800,1600)
     EXCEPT(IORATE(IOSRV(0030),LE,100))
(1)   EXCEPT(IORATE(DVUTL(6D8),GE,3))
     EXRPTS(IORATE(CHAN,DEVICE(DASD)))
(2)   EXCEPT((CHPBSY(01),GE,10))
(3)   EXCEPT(USERWORK(CPUBSY(0),LE,80))
     EXCEPT(USERWORK(AVGTSO,GE,10))
```

Post Processor Reports

RMF provides the following types of printed reports: interval reports, duration reports, summary reports, exception reports, and plot reports. Interval reports can be printed during an RMF session, or the post processor can print them at another time, for example, during off-shift hours. Samples of the interval reports printed either during a session or by the post processor appear in Chapters 3 and 4 ("Monitor I Session Interval Reports" and "Monitor II Session Reports").

Post processor reports offer a choice of reporting methods that produce overview information about your system's performance. The post processor reports are:

- Duration reports, which summarize a particular system activity over a specified reporting period. See "Duration Reports" on page 5-41 later in this chapter.
- Summary reports, which present an overview of system activity over a specified reporting period. See "Summary Reports" on page 5-44 later in this chapter.
- Exception reports, which present a summary of the values that exceeded installation-defined thresholds over a specified period of time. See "Exception Reports" on page 5-63 later in this chapter.
- Plot reports, which present a graphic summary of a number of key system performance indicators. See "Plot Reports" on page 5-51 later in this chapter.

Only the post processor can produce duration reports, summary reports, exception reports, and plot reports. The post processor generates these reports from SMF records built during a Monitor I session. While the post processor can generate interval reports from data collected during a Monitor II background session and from trace data and enqueue activity data collected during a Monitor I session, this data is not used to generate the duration, summary, exception, or plot reports.

The post processor generates interval reports from SMF records built during a single interval; however, the other post processor reports use data in SMF records that were built over a **reporting period**. The reporting period, though it could be equivalent to a single Monitor I session measurement interval, would normally span multiple intervals. The reporting period need not be a single block of continuous time; it could consist of multiple non-continuous time blocks. You could, for example, request a summary report that covers a particular time period -- such as a time of very heavy system use -- over a span of several days. Thus, the reporting period is the time span covered in a post processor report; the ranges of times and days that make up the reporting period are defined on post processor control statements, which are described earlier in this chapter.

Duration Reports

The post processor can produce duration reports for the following areas of system activity, all of which are measured during a Monitor I session:

- Workload activity
- Processor (CPU) activity
- Page/swap data set activity
- Device activity
- Paging activity
- Channel activity
- I/O queuing activity
- Virtual storage activity

Duration reports are particularly useful when you are using an RMF Monitor I session as a high-level measurement tool. In this situation, you might run the session over long periods of time and specify SMF record output. You could then use the post processor to generate duration reports. In this way, you can measure your system's performance over long periods of time with a minimal amount of system overhead and a minimal volume of printed output.

A duration report is similar in format and content to the interval report for the same system activity. However, the duration report summarizes that system activity during all of the RMF measurement intervals that fall within the duration interval. The **duration interval** is the period of time covered in the duration report. Factors to consider in selecting the length of the duration interval and its relationship to the length of the post processor reporting period are given under "DINTV Control Statement" earlier in this chapter. In any case, the recommended maximum length of a duration interval is 24 hours. The amount of data gathered during a duration interval that exceeds 24 hours could cause some or all of the data fields in the report to overflow. When this condition occurs, the field that overflows contains either a negative number or asterisks, an indication that the length of the duration interval should be reduced.

The data fields in the duration report are identical to those in the corresponding interval report, with the following minor differences:

- In the processor activity report, any processor that moved online or offline during the duration interval is indicated by an asterisk following the processor identifier.
- In the channel activity report, any channel that moved online or offline during the duration interval is indicated by an asterisk following the channel identifier.
- In the I/O device activity report, any device that moved online or offline during the duration interval is indicated by an asterisk following the device number.
- In the device activity report, if any devices are changed or deleted from a storage group, RMF replaces the name of the storage group with **CHGD** in the STORAGE GROUP column.
- In the workload activity report, a warning message appears in the report whenever the IPS changed during the duration interval; the message indicates that the data might not be consistent.

Each of these differences is noted in Chapter 3 under "Monitor I Session Reports."

In a PR/SM environment in LPAR mode, RMF provides another section of the CPU activity duration report. The partition data report section provides information on the configured partitions. The post processor accumulates only similar SMF record types when the CPU activity report is requested. The first record determines the type of records to be accumulated. For example, if the first SMF record RMF encounters is a PR/SM SMF record, RMF accumulates only PR/SM SMF records. Non-PR/SM SMF records are skipped. See Chapter 7 for more information on SMF records. *System Management Facilities* describes each SMF record RMF produces.

The following hierarchy exists when the post processor encounters SMF records that RMF writes while running in different PR/SM environments:

1. If the SMF records contain different system identifiers then the records are processed separately.

   For example, if SMF records written in partition 1 and partition 2 have different system identifiers, and the SYSID control statement is not used, then two separate reports will be generated. If the SYSID control statement is specified, only the system identified in the statement will be reported.

2. If the SMF records contain the same system identifiers but differences in MVS partition name, number of configured partitions or partition name, the records are processed selectively.

   The first SMF record encountered defines the type of records to be accumulated. Any subsequent records that do not have the same characteristics are skipped.

3. If the SMF records contain the same system identifier but differences in number of physical processors, status, wait completion or number of logical processors, the records are processed as if they were from the same system. All records are processed. When a new value is encountered an asterisk (*) appears next to the changed value on the report. The new value is reported.

The heading fields in the duration report are also similar to those in the corresponding interval report. The heading fields identify the type of operating system, the release number and level of the operating system, the four-character SMF identifier associated with the system at system generation, and the RMF report the date, in the form mm/dd/yy, and the time, in the form hh.mm.ss, when the first RMF measurement interval within the duration interval began. The END field shows the date and time when the last interval began. The INTERVAL field shows the total length of the RMF measurement intervals included in the report, in the form of hhh.mm.ss.

The CYCLE heading field indicates the length of the sampling cycle in effect during the reporting period. To set the contents of this field, the post processor uses the cycle length contained in the first SMF record it encounters. If any subsequent SMF records indicate a conflicting cycle length, all such records are ignored, and RMF issues a message that describes the records omitted from the report.

When generating reports with the post processor, take note of the version and release of RMF that you are using. An abend will occur when a non-PR/SM level of the post processor is used to process SMF records that were written by an RMF system with PR/SM support.

Figure 5-13 shows a sample duration report.

C P U A C T I V I T Y

PAGE 1

OS/VS2	SYSTEM ID AQFT	DATE 03/22/89	INTERVAL 29.59.957
SP2.2.3	RPT VERSION 3.5.1	TIME 10.00.00	CYCLE 1.000 SECONDS

CPU MODEL 3090 VERSION 61

CPU NUMBER	VF ONLINE	VF AFFINITY PERCENTAGE	BUSY TIME PERCENTAGE	WAIT TIME PERCENTAGE	CPU SERIAL NUMBER	I/O TOTAL INTERRUPT RATE	% I/O INTERRUPTS HANDLED VIA TPI
0	---	****	96.19	----	010028	29.30	9.86
1	---	****	96.11	----	110028	27.24	11.14
2	---	****	96.00	----	210028	1699	17.03
TOTAL/AVERAGE		****	96.10	0.00		1756	16.81

SYSTEM ADDRESS SPACE ANALYSIS SAMPLES = 1,800

	NUMBER OF ASIDS			DISTRIBUTION OF QUEUE LENGTHS (%)											
TYPE	MIN	MAX	AVG	0	1	2	3	4	5	6	7-8	9-10	11-12	13-14	14+
IN READY	0	33	10.8	0.2	0.6	2.0	3.7	4.5	5.5	6.2	15.0	14.0	12.0	10.8	25.1

TYPE	MIN	MAX	AVG	0	1-2	3-4	5-6	7-8	9-10	11-15	16-20	21-25	26-30	31-35	35+
IN	56	75	64.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0
OUT READY	0	4	0.1	92.1	7.1	0.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
OUT WAIT	102	137	121.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0
LOGICAL OUT RDY	0	9	0.1	90.1	8.6	0.6	0.3	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
LOGICAL OUT WAIT	19	48	33.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	7.7	21.0	40.0	30.9
BATCH	10	24	18.3	0.0	0.0	0.0	0.0	0.0	0.8	20.8	44.3	33.8	0.0	0.0	0.0
STC	61	79	67.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0
TSO	130	137	133.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0

Figure 5-13. Duration Report

Summary Reports

Summary reports provide a high-level view of system activity. A summary report can consist of one line of data that summarizes system activity for each interval within the reporting period, a single line of data giving summary totals for all RMF intervals within the reporting period, or both interval summary data lines and a total summary data line. You specify the type of reporting required on post processor control statements.

The data shown in the report is derived from the SMF records created by any Monitor I sessions that executed during the reporting period. When a particular system activity is not measured during the reporting period, the column or columns derived from measurement of that activity are omitted. For example, if a Monitor I session did not measure paging activity, the columns in the summary report that describe paging activity (SWAP RATE and DEMAND PAGING) are omitted.

The heading fields for a summary report identify the type of operating system, the release number and level of the operating system, the four-character SMF identifier associated with the system at system generation, and the RMF report level. The START field shows the date and time when the first interval in the reporting period began. The END field shows the date and time when the last interval began. The date is in the form mm/dd/yy, and the time is in the form hh.mm.ss. The INTERVAL field shows the average length of the RMF measurement interval during the reporting period, in the form hh.mm.ss. The CYCLE field shows the length of the sampling cycle during the reporting period. When all SMF records have the same cycle length, that value is reported. When different cycle lengths are encountered, the post processor sets the CYCLE field equal to the average of all cycle lengths encountered.

When a summary report consists of more than one page, the heading fields are repeated for each page. The START, END, CYCLE, and INTERVAL fields reflect the contents of the page on which they appear. When total summary data is requested, a total summary line is generated for the intervals covered on each page, and the last page of the report shows values for START, END, CYCLE, and INTERVAL that reflect the contents of all pages in the report.

Figure 5-14 shows a sample summary report, including both an interval summary line for each RMF measurement interval and a total summary line for all of the intervals. The data fields included in the report are described in the following section.

R M F   S U M M A R Y   R E P O R T

PAGE 1

OS/VS2 SYSTEM ID AQFT START 03/04/87-00.00.01 INTERVAL 00.27.42
SP2.2.0 RPT VERSION 3.5.0 END 03/04/87-23.00.00 CYCLE 1.000 SECONDS

NUMBER OF INTERVALS 32 TOTAL LENGTH OF INTERVALS 14.46.46

DATE MM/DD	TIME HH.MM.SS	INT MM.SS	CPU BUSY	DASD RESP	DASD RATE	TAPE RATE	JOB MAX	JOB AVE	TSO MAX	TSO AVE	STC MAX	STC AVE	SWAP RATE	DEMAND PAGING	SERVICE RATE	TRANS RATE
3/04	00.00.01	29.58	4.4	23	102.3	57.8	15	9	9	8	55	54	0.00	0.13	5196	0.075
3/04	00.30.00	29.59	4.4	25	102.3	61.7	12	9	7	6	55	54	0.00	0.13	4537	0.050
3/04	01.00.00	30.00	4.0	31	81.8	62.8	12	10	6	5	54	53	0.00	0.58	3729	0.141
3/04	01.30.00	29.59	5.4	36	64.8	66.1	12	8	5	4	58	52	0.00	1.13	6986	0.268
3/04	02.00.00	05.30	1.7	32	12.0	0.0	1	0	3	2	23	22	0.00	0.52	655	0.048
3/04	02.21.27	08.32	12.3	52	120.9	0.0	7	3	1	0	59	41	0.00	0.48	15192	0.144
3/04	02.30.00	30.00	9.7	52	89.6	0.0	8	3	1	0	59	37	0.00	0.01	13096	0.011
3/04	03.00.00	29.59	13.7	43	289.3	0.0	8	8	1	0	57	56	0.00	0.20	17275	0.013
3/04	03.30.00	30.00	13.7	49	299.7	0.0	9	8	0	0	63	56	0.00	0.30	16919	0.016
3/04	04.00.00	29.59	1.3	67	8.3	0.0	1	0	1	1	64	62	0.00	0.05	397	0.008
3/04	04.30.00	29.59	1.2	75	4.2	0.0	0	0	2	1	59	59	0.00	0.00	274	0.001
3/04	05.00.00	30.00	1.2	56	6.3	0.0	1	0	1	1	59	59	0.00	0.01	391	0.002
3/04	05.30.00	29.59	1.6	59	17.5	17.3	1	1	4	3	59	58	0.00	0.08	711	0.031
3/04	06.00.00	30.00	1.6	41	26.8	4.2	1	1	6	4	59	58	0.00	0.13	706	0.096
3/04	06.30.00	30.00	2.2	43	18.4	2.1	1	0	8	5	59	59	0.00	0.35	941	0.158
3/04	07.00.00	29.59	14.2	11	88.7	0.0	7	0	41	24	60	59	0.00	1.03	10271	1.317
3/04	07.30.00	30.00	19.5	10	339.2	0.0	9	3	111	71	59	56	0.00	2.54	16410	3.301
3/04	08.00.00	29.59	27.5	11	483.4	0.0	8	3	228	170	59	57	0.00	1.92	25670	6.561
3/04	08.30.00	30.00	51.9	12	792.6	0.1	15	7	297	269	58	53	0.09	2.39	50974	9.917
3/04	09.00.00	29.59	48.1	12	700.4	14.2	13	6	344	325	57	54	0.21	2.04	51795	10.500
3/04	09.30.00	30.00	61.6	12	805.5	11.1	16	7	373	356	58	53	0.38	2.13	63817	12.000
3/04	10.00.00	30.00	67.5	12	679.8	3.1	16	8	386	379	57	53	0.41	2.48	72932	11.847
3/04	10.30.00	30.00	72.4	13	898.2	7.3	16	10	412	394	54	50	0.36	2.15	110840	12.701
TOTAL/AVERAGE			28.9	18	396.8	12.4	22	7	412	156	64	53	0.14	1.20	44159	4.786

Figure 5-14. Summary Report

Data Fields

Two fields precede the data fields. NUMBER OF INTERVALS indicates the number of RMF measurement intervals included in the reporting period. TOTAL LENGTH OF INTERVALS indicates (in the form hh.mm.ss) the total length of the reporting period.

When a summary report consists of more than one page, the NUMBER OF INTERVALS field and the TOTAL LENGTH OF INTERVALS field reflect the contents of the page on which they appear. When total summary data is requested, a total summary line is generated for the intervals covered on each page, and the last page of the report shows values for NUMBER OF INTERVALS and TOTAL LENGTH OF INTERVALS and presents a total summary data line that reflect the contents of all pages in the report.

The meaning of each field in the report is described in the following text. The first three fields, which describe the date, starting time, and length of each RMF measurement interval included in the reporting period, appear only when you request interval summary data. All other fields appear for both interval summary data and total summary data. RMF omits a field (other than date, starting time, and interval time) if all values within the column are zero during the reporting period. The single line that reports total summary data begins with the heading TOTAL/AVERAGE.

DATE
MM/DD
The date, in the form mm/dd, when each interval included in the summary report began. This field is reported when you specify interval summary data; it does not appear in the total summary data line.

TIME
HH.MM.SS
The start time for each interval included in the summary report, in the form hh.mm.ss. This field is reported when you specify interval summary data; it does not appear in the total summary line.

INT MM.SS
MM.SS
The actual length of each interval included in the summary report, in the form mm.ss. This field is reported when you specify interval summary data; it does not appear in the total summary line.

CPU
BUSY
The average busy percentage during the reporting period for all processors.

In the total summary line, this field indicates the average processor-busy percentage for all processors over all of the measurement intervals included in the reporting period (or page).

This column is omitted when processor (CPU) activity was not measured during the reporting period.

**DASD
RESP**

The average number of milliseconds required to complete an I/O request on all direct access storage devices included in the report.

The rate is calculated by adding the average active time for all direct access devices to the average subchannel queuing time for all direct access devices.

In the total summary line, this field indicates the average response time for all direct access devices and all of the measurement intervals included in the reporting period (or page).

This column is omitted when I/O device activity for direct access storage devices was not measured during the reporting period.

**DASD
RATE**

The activity per second for all direct access storage devices included in the report. The value reported corresponds to an accumulation of each DEVICE ACTIVITY RATE field in the direct access device activity report.

The rate is calculated by dividing the number of successful SSCH instructions issued to the devices during the interval by the length of the interval. The range is from 0.0 to 999.9.

In the total summary line, this field indicates the average activity rate for all direct access storage devices over all of the measurement intervals included in the reporting period (or page). To calculate the rate, the post processor computes the total number of successful SSCH instructions issued during all intervals included in the reporting period and divides that total by the length of all the intervals included in the reporting period.

This column is omitted when I/O device activity for direct access storage devices was not measured during the reporting period.

**TAPE
RATE**

The activity per second for all magnetic tape devices included in the report. The value reported corresponds to an accumulation of each DEVICE ACTIVITY RATE field in the magnetic tape device activity report.

The rate is calculated by dividing the number of successful SSCH instructions issued to the devices during the interval by the length of the interval. The range is from 0.0 to 999.9.

In the total summary line, this field indicates the average activity rate for all magnetic tape devices over all of the measurement intervals included in the reporting period (or page). To calculate the rate, the post processor computes the total number of successful SSCH instructions issued during all intervals included in the reporting period and divides that total by the length of all intervals included in the reporting period.

This column is omitted when I/O device activity for magnetic tape devices was not measured during the reporting period.

JOB
MAX

The maximum number of batch jobs that were active during each measurement interval. The range is from 0 to 999. The value reported corresponds to the MAX number of BATCH address spaces in the processor (CPU) activity interval report.

In the total summary line, this field indicates the maximum number of batch jobs noted during any of the measurement intervals included in the reporting period (or page).

This column is omitted when processor (CPU) activity was not measured during the reporting period.

JOB
AVG

The average number of batch jobs that were active during each measurement interval. The range is from 0 to 999. The value reported corresponds to the AVG number of BATCH address spaces in the processor (CPU) activity interval report.

In the total summary line, this field indicates the average number of batch jobs for the reporting period (or page). The value is calculated by dividing the total number of active jobs observed by the total number of samples.

This column is omitted when processor (CPU) activity was not measured during the reporting period.

TSO
MAX

The maximum number of TSO sessions that were active during each measurement interval. The range is from 0 to 999. The value reported corresponds to the MAX number of TSO address spaces in the processor (CPU) activity report.

In the total summary line, this field indicates the maximum number of TSO sessions noted during any of the measurement intervals included in the reporting period (or page).

This column is omitted when processor (CPU) activity was not measured or TSO was not active during the reporting period.

TSO
AVE

The average number of TSO sessions that were active during each measurement interval. The range is from 0 to 999. The value reported corresponds to the AVG number of TSO address spaces in the processor (CPU) activity interval report.

In the total summary line, this field indicates the average number of TSO sessions for the reporting period (or page). The value is calculated by dividing the total number of active TSO sessions observed by the total number of samples.

This column is omitted when processor (CPU) activity was not measured or TSO was not active during the reporting period.

**STC
MAX**

The maximum number of started tasks and mount tasks that were active during each measurement interval. The range is from 0 to 999. The value reported corresponds to the MAX number of STC address spaces in the processor (CPU) activity interval report.

In the total summary line, this field indicates the maximum number of started tasks and mount tasks noted during any of the measurement intervals included in the reporting period (or page).

This column is omitted when processor (CPU) activity was not measured during the reporting period.

**STC
AVE**

The average number of started tasks and mount tasks that were active during each measurement interval. The range is from 0 to 999. The value reported corresponds to the AVG number of STC address spaces in the processor (CPU) activity interval report.

In the total summary line, this field indicates the average number of started tasks and mount tasks for the reporting period (or page). The value is calculated by dividing the total number of started tasks and mount tasks observed by the total number of samples.

This column is omitted when processor (CPU) activity was not measured during the reporting period.

**SWAP
RATE**

The number of swaps per second for each interval. The value reported corresponds to the SWAPS PER SECOND field under SWAP SEQUENCE COUNTS in the Monitor I paging activity interval report.

The rate is calculated by dividing the total number of swaps occurring during the interval by the length of the interval. The range is from 0.00 to 9.99.

In the total summary line, this field indicates the average swap rate over all of the measurement intervals included in the reporting period (or page). To calculate the rate, the post processor computes the total number of swaps during all intervals included in the reporting period and divides that total by the length of all intervals included in the reporting period.

This column is omitted when paging activity was not measured during the reporting period.

**DEMAND
PAGING**

The number of demand paging requests per second for each interval. The value reported corresponds to the sum of the PAGE RECLAIMS/NON VIO and PAGE IN/NON VIO/NON SWAP fields in the Monitor I paging activity interval report.

The rate is calculated by dividing the total number of non-VIO, non-swap page-ins and page reclaims by the length of the interval. The range is from 0.00 to 999.99.

The demand paging rate is the same as the page fault rate described in the SRM section of the *Initialization and Tuning*. Although the values are derived in different ways, they are both indicators of contention for real storage resources.

In the total summary line, this field indicates the average demand paging rate over all of the measurement intervals included in the reporting period (or page). To calculate the rate, the post processor computes the total number of non-VIO, non-swap page-ins and page reclaims during all intervals included in the reporting period. That total is then divided by the length of all intervals included in the reporting period.

This column is omitted when paging activity was not measured during the reporting period.

SERVICE RATE

The number of service units (input/output service, processor service, and main storage service) absorbed during each second of the interval. The value reported corresponds to the INTERVAL SERVICE/PER SEC= field in the system summary workload activity interval report.

The rate is calculated by dividing the total number of service units absorbed by the length of the interval. The range is from 0 to 9,999,999.

In the total summary line, this field contains the number of service units absorbed per second over the total length of the reporting period (or page). To calculate the rate, the post processor computes the total number of service units absorbed during all intervals included in the reporting period and divides that total by the length of all intervals included in the reporting period.

This column is omitted when workload activity was not measured during the reporting period.

Note: When the SET command is used to change the IPS during an interval, an asterisk (*) appears after the value for TRANS RATE. The asterisk indicates that the value reported for both TRANS RATE and SERVICE RATE reflects only the part of the interval that had elapsed before the SET command. A SET command that changes the installation control specification has the same effect.

TRANS RATE

The average number of transactions that ended during each second of the interval. The value reported comes from the same source as the ENDED TRANSACTIONS field in the workload activity interval report.

The rate is calculated by dividing the total number of ended transactions for the interval by the length of the interval. The range is from 0.000 to 9.999.

In the total summary line, this field contains the average rate of ended transactions over the reporting period (or page). To calculate the rate, the post processor computes the total number of ended transactions during all intervals in the reporting period and divides that total by the length of all intervals included in the reporting period.

This column is omitted when workload activity was not measured during the reporting period.

Note: When the SET command is used to change the IPS during an interval, an asterisk (*) appears after the value for TRANS RATE. The asterisk indicates that the value reported for both TRANS RATE and SERVICE RATE reflects only the part of the interval that had elapsed before the SET command. A SET command that changes the installation control specification has the same effect.

Plot Reports

Plot reports present a graphic summary of system activity over a specified time range -- the plot reporting period. Plot reports, like the other post processor reports, are generated from data contained in SMF records built during a Monitor I session. The post processor selects the SMF records to be used according to the type of plots that you request and the reporting period that you specify. You describe the type of plot reporting you require on control statements for the post processor. These control statements are:

- PLOTS, which defines the types of plots you require. See "PLOTS Control Statement" earlier in this chapter for details. The description of each different type of plot later in this section includes the option you must specify on the PLOTS statement for each report.
- PINTV, which specifies the length of the plot interval. The **plot interval** is the length of time that is to elapse between plotting each line in the report. For most effective plot reporting, the plot interval should be either equal to the RMF measurement interval or a multiple of the RMF measurement interval. See "PINTV Control Statement" earlier in this chapter for more details.
- PTOD, which specifies the length of time to be included in the plot report for each day in the reporting period. See "PTOD Control Statement" earlier in this chapter.
- MAXPLEN, which specifies the maximum number of lines to be included in each plot report. See "MAXPLEN Control Statement" earlier in this chapter.

Each page of plot report output consists of headings and space for two plot reports. The heading fields identify the type of operating system, the release number and level of the operating system, the four-character SMF identifier associated with the system at system generation, and the RMF report level. The START field shows the date, in the form mm/dd/yy, and the time, in the form hh.mm.ss, when the first RMF measurement interval in the plot reporting period began. The END field shows the date and time when the last interval began. The INTERVAL field shows the length of the plot interval -- the value specified on the PINTV control statement or, when the PINTV statement was omitted, the average length of the RMF Monitor I measurement interval during the reporting period, in the form hh.mm.ss. The CYCLE field shows the average length of the cycle at which data was sampled during the plot reporting period.

Each plot report consists of a heading that identifies the plot report, the measurement scale for that report, and a line of plot data for each plot interval within the plot reporting period.

The measurement scale can be either a percentage range or a range of fixed numbers, depending on the type of plot that you request. It is possible that the measurement scale provided for each plot cannot reflect the activity of your particular system accurately; perhaps the range is too small or too large. When the value recorded for a plot line exceeds the upper limit of the measurement scale, RMF prints plus signs (+) across the entire line. A field within the line of plus signs indicates the value recorded. Should your reports indicate that a measurement scale is not adequate for your particular installation, you can modify the measurement scale for any type of plot. The procedure to follow and an example of how to modify the measurement scale for a plot report appear later in this chapter under "Modifying the Measurement Scale."

Figure 5-15 shows a sample of plot report output. As the figure shows, each line in the plot report consists of a line of asterisks that indicate the point on the measurement scale recorded for each plot interval. Each line in the plot is preceded by the time when the RMF measurement interval that collected the plotted data began; the first plot line in each calendar day included in the report is indicated by the date, in the form mm/dd. When more than one RMF interval is used for each plot line, the data from the intervals used is combined. The time shown is the time for the first measurement interval.

Types of Plot Reports

The following areas of system activity can be reported in the form of a plot. Each plot report description includes:

- The name of the plot report.
- The option used to request the plot report.
- The system activity measured in the plot report.
- The heading and measurement scale for the plot report.
- The type of Monitor I session measurement activity that is required to generate the plot report.

The plot reports are described in alphabetical order.

```
                                                  R M F   P L O T   R E P O R T
            OS/VS2                   SYSTEM ID AQFT               START 03/03/87 23.30.00  INTERVAL 00.27.29
            SP2.2.0                  RPT VERSION 3.5.0            END   03/05/87 15.30.00  CYCLE 1.000 SECONDS

            CPU BUSY PERCENTAGE - CPU 0                                        CPU BUSY PERCENTAGE - CPU 1
           0      20.0      40.0      60.0      80.0     100.0           0      20.0      40.0      60.0      80.0     100.0
           |---------|---------|---------|---------|---------|           |---------|---------|---------|---------|---------|
03/03 23:30|******** .         .         .         .         . 03/03 23:30|*********.         .         .         .         .
03/04 00:00|*****    .         .         .         .         . 03/04 00:00|******   .         .         .         .         .
      00:30|*****    .         .         .         .         .       00:30|******   .         .         .         .         .
      01:00|*****    .         .         .         .         .       01:00|******   .         .         .         .         .
      01:30|******   .         .         .         .         .       01:30|******** .         .         .         .         .
      02:00|**       .         .         .         .         .       02:00|***      .         .         .         .         .
      02:21|**********         .         .         .         .       02:21|**********         .         .         .         .
      02:30|*********.         .         .         .         .       02:30|******** .         .         .         .         .
      03:00|***************    .         .         .         .       03:00|****************   .         .         .         .
      03:30|**************     .         .         .         .       03:30|***************    .         .         .         .
      04:00|**       .         .         .         .         .       04:00|**       .         .         .         .         .
      04:30|**       .         .         .         .         .       04:30|**       .         .         .         .         .
      05:00|**       .         .         .         .         .       05:00|**       .         .         .         .         .
      05:30|**       .         .         .         .         .       05:30|**       .         .         .         .         .
      06:00|**       .         .         .         .         .       06:00|***      .         .         .         .         .
      06:30|**       .         .         .         .         .       06:30|***      .         .         .         .         .
      07:00|***********        .         .         .         .       07:00|************       .         .         .         .
      07:30|****************   .         .         .         .       07:30|****************** .         .         .         .
      08:00|**********************       .         .         .       08:00|***********************      .         .         .
      08:30|********************************       .         .       08:30|*********************************      .         .
      09:00|*******************************        .         .       09:00|*******************************        .         .
      09:30|***********************************    .         .       09:30|************************************   .         .
```

Figure 5-15. Plot Report

Batch Users (BATCH): plots the maximum number of batch users. The value reported corresponds to the MAX number of BATCH address spaces in the processor (CPU) activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MAXIMUM BATCH UTILIZATION
0         20        40        60        80        100
|---------|---------|---------|---------|---------|
```

The maximum number of batch users can be plotted only when processor (CPU) activity was measured during the reporting period.

Channel Path Busy (CH): plots the percentage of time when a channel path was busy. The value reported corresponds to the PERCENT CH PATH BUSY field in the channel path activity interval report.

The measurement scale consists of percentages. The heading and measurement scale for the report are:

```
 CHANNEL PATH nn         BUSY PERCENTAGE
0      10.0      20.0      30.0      40.0      50.0
|---------|---------|---------|---------|---------|
```

CHANNEL PATH nn is the hexadecimal channel path identifier. The channel path identifier uniquely identifies the channel path and that must be specified if this plot report is requested.

The channel path busy percentage can be plotted only when channel path activity was measured during the reporting period.

CSA Allocated (CSAA): plots the maximum amount of allocated CSA below the 16-megabyte line. The value reported corresponds to the MAX amount for ALLOCATED CSA in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
MAXIMUM ALLOCATED CSA (IN K BYTES)
0        500      1000      1500      2000     2500
|---------|---------|---------|---------|---------|
```

Allocated CSA can be plotted only when virtual storage activity was measured during the reporting period.

CSA Free (CSAFP): plots the minimum number of free CSA pages below the 16-megabyte line. The value reported corresponds to the MIN amount for CSA FREE PAGES (BYTES) in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
MINIMUM CSA FREE PAGES (IN K BYTES)
0        500      1000      1500      2000      2500
|---------|---------|---------|---------|---------|
```

Free CSA can be plotted only when virtual storage activity was measured during the reporting period.

Control Unit Busy Delay Time (CUBDL): plots the average number of milliseconds when any I/O requests for the device were delayed because the control unit was busy. The value corresponds to the AVG CUB DELAY field in the Monitor I device activity report. RMF can plot the control unit delay time only when the device activity was measured and RMF was running on a 3090 processor during the reporting period.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEVICE nnn    AVERAGE CONTROL UNIT BUSY DELAY TIME
0      20.0      40.0      60.0      80.0     100.0
|---------|---------|---------|---------|---------|
```

DEVICE **nnn** identifies the device included in the report; **nnn** is a three-digit hexadecimal number that uniquely identifies the device and that must be specified if the plot report is requested.

Demand Paging Rate (PAGE): plots the demand paging rate. The rate reported is the number of non-VIO, non-swap page-ins and page reclaims per second. The rate corresponds to the sum of the PAGE RECLAIMS/NON VIO and PAGE IN/NON VIO/NON SWAP fields in the Monitor I paging activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEMAND PAGING RATE
0      40.0      80.0     120.0     160.0     200.0
|---------|---------|---------|---------|---------|
```

The demand paging rate can be plotted only when paging activity was measured during the reporting period.

Device Activity Rate (DAC): plots the device activity rate -- the number of successful SSCH instructions per second issued to the device -- for a specific device. For a multiple exposure device, the activity rate is the sum of the activity rates for all exposures. The rate corresponds to the DEVICE ACTIVITY RATE field in the I/O device activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEVICE nnn               ACTIVITY RATE
0     3.000     6.000     9.000    12.000    15.000
|---------|---------|---------|---------|---------|
```

DEVICE **nnn** identifies the device included in the report; **nnn** is a three-digit hexadecimal device number that uniquely identifies the device and that must be specified of the plot report is requested.

An asterisk preceding the time of a plot data line indicates that records were ignored because the sampling cycle changed during the reporting period.

The device activity rate can be plotted only when I/O device activity for the appropriate device class was measured during the reporting period.

Device Busy Delay Time (DBDL): plots the average number of milliseconds when any I/O requests for the device were delayed because the device was busy. The value corresponds to the AVG DB DELAY field in the Monitor I device activity report. RMF can plot the control unit delay time only when the device activity was measured and RMF was running on a 3090 processor during the reporting period.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEVICE nnn    AVERAGE DEVICE BUSY DELAY TIME
0       20.0      40.0      60.0      80.0     100.0
|---------|---------|---------|---------|---------|
```

DEVICE **nnn** identifies the device included in the report; **nnn** is a three-digit hexadecimal number that uniquely identifies the device and that must be specified if the plot report is requested.

Device Connect Time (DCON): plots the average number of milliseconds per request the device was connected to a channel path.

The measurement scale consists of fixed numbers. the heading and measurement scale for the report are:

```
 DEVICE nnn             AVERAGE CONNECT TIME
0        20.0      40.0      60.0      80.0     100.0
|---------|---------|---------|---------|---------|
```

DEVICE nnn identifies the device included in the report; nnn is a three-digit hexadecimal device number that uniquely identifies the device and that must be specified if the plot report is requested.

An asterisk preceding the time of a plot data line indicates that records were ignored because the sampling cycle changed during the reporting period.

The device connect time can be plotted only when I/O device activity for the device was measured during the reporting period.

Device Response Time (DRT): plots the average number of milliseconds required to complete an I/O request.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEVICE nnn          AVERAGE RESPONSE TIME
0       20.0      40.0      60.0      80.0     100.0
|---------|---------|---------|---------|---------|
```

DEVICE nnn identifies the device included in the report; nnn is a three-digit hexadecimal device number that uniquely identifies the device and that must be specified if the plot report is requested.

An asterisk preceding the time of a plot data line indicates that records were ignored because the sampling cycle changed during the reporting period.

The device response time can be plotted only when I/O device activity for the appropriate device class was measured during the reporting period.

Extended Storage Migration Rate (ESMR): plots the average rate of page migration from extended to auxiliary storage. The rate corresponds to the MIGRATION RATE in the Monitor I paging activity report. RMF can plot the migration rate only when paging activity was measured during the reporting period and the system has installed extended storage.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MIGRATION RATE FROM EXTENDED TO AUXILIARY STORAGE
0       40.0      80.0     120.0     160.0     200.0
|---------|---------|---------|---------|---------|
```

Average HUIC Count (HUIC): plots the average high unreferenced interval count. The value corresponds to the AVG HIGH UIC field in the Monitor I paging activity report. RMF can plot the average high UIC only when paging activity was measured during the reporting period.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 AVERAGE HIGH UIC COUNT
0        60.0     120.0     180.0     240.0     300.0
|---------|---------|---------|---------|---------|
```

I/O Activity Rate (IOAC): plots the I/O activity rate -- the number of requests that were successfully selected from the logical control unit queue per second. For 308x, 908x, and 4381 processors, the rate corresponds to the ACTIVITY RATE field in the I/O queueing activity report.

For the 3090 processor, the activity rate is calculated by adding the path taken count for each path of the control unit per second. There is no corresponding field in the Monitor I session I/O queuing activity report. In that report, the field CHPID TAKEN shows the separate rates for each channel path.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
   LCU nnn            ACTIVITY RATE
0     50.000     100.000   150.000   200.000  250.000
|---------|---------|---------|---------|---------|
```

LCU **nnn** identifies the logical control unit included in the report; **nnn** is a one, two or three-digit hexadecimal number that uniquely identifies the logical control unit and that must be specified if this plot report is requested. The LCU number can be specified as a one to three-digit number.

An asterisk preceding the time of a plot data line indicates that records were ignored because the sampling cycle changed during the reporting period.

The activity rate can be plotted only when I/O queueing activity was measured during the reporting period.

Page Placement on Extended Storage (PTES): plots the average paging transfer rate to extended storage. The rate corresponds to the RATE OF PAGE MOVEMENT TO ES field in the Monitor I paging activity report. RMF can plot the page transfer rate to extended storage only when paging activity was measured during the reporting period and the system has installed extended storage.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 PAGE PLACEMENT ON EXTENDED STORAGE
0      80.0      160.0     240.0     320.0     400.0
|---------|---------|---------|---------|---------|
```

LCU **nnn** identifies the logical control unit included in the report; **nnn** is a two-digit hexadecimal number that uniquely identifies the logical control unit and that must be specified if this plot report is requested. The LCU number can be specified as a one to three-digit number.

Processor Busy (CPUID): plots the percentage of time when the processor was busy. The percentage comes from the same source as the WAIT TIME PERCENTAGE field in the processor (CPU) activity interval report.

The value reported is calculated by subtracting the wait time percentage for a processor from 100 percent.

The measurement scale consists of percentages. The heading and measurement scale for the report are:

```
 CPU BUSY PERCENTAGE - CPU n
0      20.0      40.0      60.0      80.0     100.0
|---------|---------|---------|---------|---------|
```

CPU **n** identifies the processor included in the report, where **n** is the processor identifier.

The percentage of time a processor was busy can be plotted only when processor (CPU) activity was measured during the reporting period.

Service Rate (SEV): plots the service rate -- the number of service units provided per second. The rate corresponds to the INTERVAL SERVICE/PER SEC= field in the workload activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 SERVICE RATE
0      2000      4000      6000      8000     10000
|---------|---------|---------|---------|---------|
```

An asterisk preceding the time of a plot data line indicates that the IPS changed during the reporting period.

The service rate can be plotted only when workload activity was measured during the reporting period.

SQA Allocated (SQAA): plots the maximum amount of allocated SQA below the 16-megabyte line. The value reported corresponds to the MAX amount for ALLOCATED SQA in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
MAXIMUM ALLOCATED SQA (IN K BYTES)
0         500       1000      1500      2000      2500
|---------|---------|---------|---------|---------|
```

Allocated SQA can be plotted only when virtual storage activity was measured during the reporting period.

SQA Expansion (SQAE): plots the amount of SQA expansion into CSA below the 16-megabyte line. The value reported corresponds to the SQA EXPANSION INTO CSA field in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
AMOUNT OF SQA EXPANSION INTO CSA (IN K BYTES)
0         500       1000      1500      2000      2500
|---------|---------|---------|---------|---------|
```

SQA expansion into CSA can be plotted only when virtual storage activity was measured during the reporting period.

SQA Free (SQAFP): plots the minimum number of free SQA pages below the 16-megabyte line. The value reported corresponds to the MIN amount for SQA FREE PAGES (BYTES) in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
MINIMUM SQA FREE PAGES (IN K BYTES)
0        500       1000      1500      2000      2500
|---------|---------|---------|---------|---------|
```

Free SQA can be plotted only when virtual storage activity was measured during the reporting period.

Started Task and Mount Task Users (STC): plots the maximum number of started task and mount task users. The value reported corresponds to the MAX number of STC address spaces in the processor (CPU) activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MAXIMUM STC UTILIZATION
0        10        20        30        40        50
|---------|---------|---------|---------|---------|
```

The number of started task and mount task users can be plotted only when processor (CPU) activity was measured during the reporting period.

Swap Rate (SWA): plots the swap sequence rate -- the number of times per second that storage was swapped out and then swapped in. The rate corresponds to the SWAPS PER SECOND field in the Monitor I paging activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 SWAP RATE
0       2.000     4.000     6.000     8.000    10.000
|---------|---------|---------|---------|---------|
```

The swap rate can be plotted only when paging activity was measured during the reporting period.

Total Paging Rate (TPAG): plots the total paging rate. To calculate the rate, the post processor determines the total number of page-ins and page-outs that occurred for any reason during each plot interval; this total is divided by the number of seconds in the plot interval. Page reclaims are not included in the calculation; the rate reported thus represents the I/O activity performed because of paging. The rate corresponds to the sum of the SUM/TOTAL RATE/PAGE IN and SUM/TOTAL RATE/PAGE OUT fields in the Monitor I paging activity report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 TOTAL PAGING RATE
0      100.0     200.0     300.0     400.0     500.0
|---------|---------|---------|---------|---------|
```

The total paging rate can be plotted only when paging activity was measured during the reporting period.

Transaction Rate (TRA): plots the number of ended transactions per second. The rate reported comes from the same source as the ENDED TRANSACTIONS field in the workload activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 TRANSACTION RATE
0        4.000     8.000    12.000    16.000    20.000
|---------|---------|---------|---------|---------|
```

An asterisk preceding the time of a plot data line indicates that the IPS changed during the reporting period.

The transaction rate can be plotted only when workload activity was measured during the reporting period.

TSO Sessions (TSO): plots the maximum number of TSO sessions. The value reported corresponds to the MAX number of TSO address spaces in the processor (CPU) activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MAXIMUM TSO UTILIZATION
0         60        120       180       240       300
|---------|---------|---------|---------|---------|
```

The number of TSO sessions can be plotted only when processor (CPU) activity was measured and TSO was active during the reporting period.

Modifying the Measurement Scale

Measurement scales are provided for each plot report; these scales reflect average system use. However, you might find that some or all of the measurement scales provided are not suitable for reflecting actual system use at your installation. Such a condition is indicated when plot lines consistently overflow the measurement scale or when plot lines consistently cluster on the low end of the measurement scale. When either condition occurs, you can modify the upper limit of the measurement scale for any plot that requires a change. When you modify the upper limit, the major break points within the scale are modified automatically.

The upper limits of each measurement scale appear in the data control section ERBMFPLD within the ERBRMFPP load module. Figure 5-16 shows each plot option, the offset into control section ERBMFPLD required to access the upper limit for each plot, the contents of each entry in ERBMFPLD, and the upper limit of each measurement scale.

Each entry consists of two words. The first word contains the maximum value. The second word contains the scale factor, that is, the power of ten by which the maximum value is multiplied to arrive at the upper limit. For example, the upper limit for the swap rate plot (SWA) is 10.000; this is computed by multiplying the maximum value (10000) by the scale factor (10ff^{8} or .001). Note that the scale factor must not be modified; you change the upper limit by modifying the maximum value only.

To change the scale for a particular report, perform the following steps:

1. Determine the upper limit that you require for the scale you are modifying. The values RMF prints when a plot line exceeds the upper limit can help you to determine a new upper limit that is suitable for your installation. The limit you choose must be a multiple of five.
2. Use Figure 5-16 to locate the offset to the maximum value for the scale you are modifying. Determine the maximum value that, multiplied by the scale factor, will yield the upper limit that you want. Convert that number to hexadecimal.
3. Run AMASPZAP to update the maximum value for the entry in ERBMFPLD that describes the plot whose measurement scale you are modifying.

PLOTS Option	Offset (hex)	ERBMFPLD Entry Maximum Value DEC	HEX	Scale Factor	Upper Unit
CPUID	0	1000	(3E8)	-1	100.0
STC	8	50	(32)	0	50
TSO	10	300	(12C)	0	300
BATCH	18	100	(64)	0	100
CH	20	500	(1F4)	-1	50.0
DCON	28	1000	(3E8)	-1	100.0
DAC	30	15000	(3A98)	-3	15.000
DRT	38	100	(64)	0	100
PAGE	40	2000	(7D0)	-1	200.0
TPAG	48	5000	(1388)	-1	500.0
SWA	50	10000	(2710)	-3	10.000
SEV	58	10000	(2710)	0	10000
TRA	60	20000	(4E20)	-3	20.000
IOAC	68	250000	(3D090)	-3	250.000
SQAA	70	2500	(9C4)	0	2500
CSAA	78	2500	(9C4)	0	2500
SQAE	80	2500	(9C4)	0	2500
SQAFP	88	2500	(9C4)	0	2500
CSAFP	90	2500	(9C4)	0	2500
DBDL	98	1000	(3E8)	-1	100.0
CUBDL	A0	1000	(3E8)	-1	100.0
PTES	A8	4000	(FA0)	-1	400.0
ESMR	B0	2000	(7D0)	-1	400.0
HUIC	B8	3000	(BB8)	-1	300.0

Figure 5-16. Measurement Scale Entries

Example

The measurement scale for the swap rate plot report runs from 0 to 10.000. Assume that you want to raise the upper limit to extend the scale from 0 to 20.000. The maximum value for SWA, located at offset X'50' from the beginning of ERBMFPLD, is 10000 (X'2710'). To achieve an upper limit of 20.000, the maximum value for SWA must be changed to 20000 (X'4E20').The superzap job required to change the measurement scale is:

```
//ZAPLIMIT  JOB     accounting information,MSGLEVEL=1
//S1        EXEC       PGM=AMASPZAP
//SYSPRINT  DD      SYSOUT=A
//SYSLIB    DD      DSN=SYS1.LINKLIB,DISP=OLD
//SYSIN     DD      *
    NAME    ERBRMFPP    ERBMFPLD
    VER 50      00002710
    REP 50      00004E20
/*
```


Exception Reports

An exception report presents a summary of the values that exceeded installation-defined thresholds over a specific period of time. Exception reports, like the other post processor reports, are generated from data contained in SMF records built during a Monitor I session. RMF compares the threshold values specified in the exception control statements with the computed value in the appropriate SMF record field. If the threshold is exceeded, RMF writes a line in the exception report. The exception control statements are:

- ETOD, which specifies the start time and end time for the exception report period. See "ETOD Control Statement" earlier in this chapter.
- EXCEPT, which specifies the conditions that determine whether an exception report is to be generated. See "EXCEPT Control Statement" earlier in this chapter.
- EXRPTS, which specifies the interval reports to be generated by the post processor. See "EXRPTS Control Statement" earlier in this chapter.

The heading fields for an exception report identify the type of operating system, the release number and level of the operating system, the four-character SMF identifier associated with the system at system generation, and the RMF report level. The START field shows the date and time when the first interval in the reporting period began. The END field shows the date and time when the last interval began. The date is in the form mm/dd/yy, and the time is in the form hh.mm.ss. The INTERVAL field shows the average length of the RMF measurement interval during the reporting period, in the form hh.mm.ss. The CYCLE field shows the length of the sampling cycle during the reporting period. When all SMF records have the same cycle length, that value is reported. When different cycle lengths are encountered, the post processor sets the CYCLE field equal to the average of all cycle lengths encountered.

When an exception report consists of more than one page, the heading fields are repeated for each page. The START, END, CYCLE, and INTERVAL fields reflect the contents of the data in the entire report.

Figure 5-17 shows a sample report. The data fields included in the report are described in the following section.

Data Fields

Two fields precede the data fields. NUMBER OF INTERVALS indicates the number of RMF measurement intervals included in the reporting period. TOTAL LENGTH OF INTERVALS indicates (in the form hh.mm.ss) the total of the reporting period.

When an exception report consists of more than one page, the NUMBER OF INTERVALS field and the TOTAL LENGTH OF INTERVALS field reflect the contents of the entire report.

The meaning of each field in the report is described in the following text.

CONTROL STMT
NAME
The control statement name, if one was specified, that the post processor uses to group control statements.

DATE
MM/DD
The date, in the form mm/dd, when the interval during which the exception occurred began.

TIME
HH.MM.SS
The start time for the interval during which the exception occurred, in the form hh.mm.ss.

INT
MM.SS
The actual length of the interval during which the exception occurred, in the form mm.ss.

EXCEPTION
THRESHOLD
The threshold value and the relational operator specified in the EXCEPT statement.

ACTUAL
VALUE
The actual value derived from the data contained in the SMF record. If RMF searches more than one resource for a value that exceeded the specified threshold level, the value is not printed. In addition, the field is not printed if all or part of a qualifier is omitted, or if a device qualified by class is specified.

Note: Because RMF processes the values differently, the actual values in this report may differ slightly from those in corresponding interval reports.

EXCEPTION
DESCRIPTION
A short description of the exception. This is the condition name as specified in the EXCEPT control statement.

NAME
VALUE
The condition-name qualifier as specified in the EXCEPT control statement.


```
                                      R M F   E X C E P T I O N   R E P O R T
                                                                                                              PAGE 1

          OS/VS2                  SYSTEM ID AQFT            START 03/03/87 23.30.00  INTERVAL 00.28.03
          SP2.2.0                 RPT VERSION 3.5.0         END   03/05/87 23.00.00  CYCLE 1.000 SECONDS

NUMBER OF INTERVALS 65          TOTAL LENGTH OF INTERVALS 30.24.12
CONTROL   DATE    TIME    INT  EXCEPTION        ACTUAL
STMT NAME MM/DD HH.MM.SS MM.SS THRESHOLD        VALUE  EXCEPTION DESCRIPTION     NAME VALUE
--------- ----- -------- ----- ------------ --------- ------------------------ ----------
DEVRPT    03/03 23.30.00 30.01 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 00.00.01 29.58 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 00.30.00 29.59 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 01.00.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 01.30.00 29.59 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 02.00.00 05.30 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 02.21.27 08.32 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 02.30.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 03.00.00 29.59 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 03.30.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 04.00.00 29.59 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 04.30.00 29.59 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 05.00.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 05.30.00 29.59 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 06.00.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 06.30.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 07.00.00 29.59 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 07.30.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 08.00.00 29.59 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 08.30.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 09.00.00 29.59 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 09.30.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 10.00.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
DEVRPT    03/04 10.30.00 30.00 GE 10                   DEVICE RESPONSE TIME     -DASD
```

Figure 5-17. Exception Report

Chapter 6. Procedures for Adding User Functions

Facilities in RMF allow you to gather and report data relevant to your installation.

During a Monitor I Session: user exits are provided to enable you to sample data at each RMF cycle, collect this data and examine system indicators at each RMF interval, format and write your own SMF records, and format and write your own reports. You can also use the RMF trace facilities to trace the contents of any SQA, fixed CSA, or nucleus field that you require.

During a Monitor II Session: defined interfaces enable you to gather and report your own data by coding your own data-gathering and data-reporting routines. RMF provides the USER option for a background session and the USER menu item for a display session. To generate one additional report, you replace module ERBGUS99 with your data gatherer and ERBRUS99 with your data reporter. Specifying USER then causes your own report to be generated. Should you want to obtain more than one user report, you must add an entry to the option list or menu list as well as supply a data-gathering and a data-reporting routine. Data gathered for your routine can be reported either during the session or during execution of the post processor.

Another user exit, applicable only to Monitor II TSO display sessions, enables your installation to verify that a terminal user is authorized to use RMF.

Monitor I Session User Reports

To gather and report data relevant to your installation during a Monitor I session, RMF provides both the EXITS option and user exits at various points during Monitor I session processing. When EXITS is specified, you can:

- Initialize for the other user exit routines
- Sample fixed CSA, SQA, or nucleus data at each RMF cycle
- Perform interval processing, for example, reduce sampled data, examine system state indicators, format SMF records to be written to the SMF data set or passed to your report writer
- Write reports during a session
- Handle termination processing for the other user exits
- Write reports during execution of the post processor.

In addition, you can use the Monitor I session tracing routines to trace the contents of a fixed SQA, CSA, or nucleus field regardless of whether or not EXITS is specified.

Guidelines

Each of the user functions is described in detail in the following sections. The following guidelines apply to Monitor I user exit routines:

- All of the user exit routines must be reenterable.
- All user-written exit routines receive control in 31-bit addressing mode.
- The routines must save registers when they receive control and restore registers when they return control. Register 13 contains the address of the register save area; register 14 contains the return address; and register 15 contains the entry address.
- One input parameter that RMF passes to each user exit routine (except the tracing routine and the post processor user exit) is the address of a two-word area reserved for the use of your routines. Because these words provide a means of communication between your exit routines, their use should be controlled by conventions agreed upon by your installation.
- RMF passes a phase parameter to each user exit routine except the sampler, the tracing routine, and the post processor user exit. This phase parameter indicates which RMF phase is invoking the user exit.

RMF provides dummy routines for all Monitor I session exits that are not used.

Caution: Because all of the user exit routines except ERBMFRUR (the report writer) run in supervisor state with a key of 0, your installation must carefully control their use. Program errors that cause an exit routine to overlay system areas could bring down the system.

Initialization for Monitor I Session User Exit Routines

The initialization user exit is ERBMFIUC. It is called at the start of a Monitor I session and whenever the Monitor I session options are modified. Use this exit to perform any initialization the other user exits require, such as building a control block structure.

ERBMFIUC is invoked by IGX00007. When the exit routine gets control, register 1 points to a three-word address list. The first address points to the two-word area reserved for use by your routines. This same two-word area is passed to all the user exit routines and can be used for communication between them. The second address points to the RMF phase parameter, a full-word field that is always X'4:', indicating that the exit is called during Monitor I session initialization. The third address points to a word that is relevant only when you are providing a routine to sample data at each cycle; one of the functions your initialization routine will perform is to put the address of the user sampler in this word. Figure 6-1 illustrates the input parameter structure.

When the initialization routine is entered, the system is in supervisor state, and all interrupts are enabled. ERBMFIUC runs in key 0.

Special initialization procedures are required when your user routines include a sampling routine to sample data at each cycle; see "Sampling Data at Each Cycle." When you have a user sampler, your initialization routine **must** do the following:

- The user sampling routine must be loaded and page fixed. You must use the PGSER macro to page fix the user sampler routine because the sampler code runs disabled.
- The address of the user sampling routine must be placed in the third input parameter.
- All storage the sampler routine will require must be obtained; this storage must be obtained from SQA (subpool 245).
- The address of the SQA storage obtained must be placed in one of the two user words. The choice depends on the conventions established at your installation.

When you have completed the initialization required by all the user exits, return control by branching on register 14.

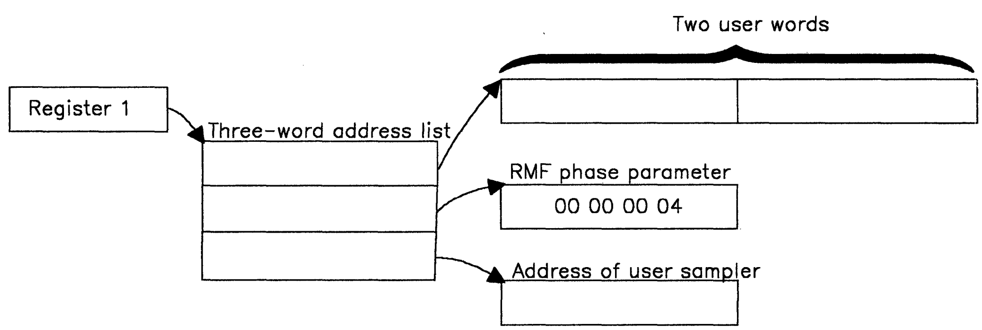

Figure 6-1. ERBMFIUC Input Parameter Structure

Sampling Data at Each Cycle

To sample data at each cycle, the steps described earlier for initialization must be performed to load and page fix the user sampler routine. A user sampler routine is activated at each cycle only when another measurement that includes a sampling routine is activated. These measurements include paging activity, page/swap data set activity, channel path activity, I/O queuing activity, device activity, and ASM/RSM/SRM tracing. At least one of these measurements must be specified to enable RMF to invoke your user sampler.

The user sampler is invoked by ERBMFEVT. When the sampler gets control, register 1 points to a two-word area. One of these words, selected by your installation, contains the address of the storage area obtained for the sampler by ERBMFIUC. Figure 6-2 illustrates the input parameter structure.

When the user sampler is entered, the system is in supervisor state, and all interrupts are disabled. The routine runs in key 0. It can sample any fixed data in CSA, SQA, or the nucleus; no other data areas can be sampled. You place the data sampled in the storage area obtained by ERBMFIUC and passed to you when your routine is invoked. This storage area is always in SQA (subpool 245). At the end of the RMF interval, RMF passes the address of the storage area to the user interval processing routine. Should your routine cause a page fault, the Monitor I session terminates abnormally with an abend code of 0FE.

When your sampling is completed, return control by branching on register 14.

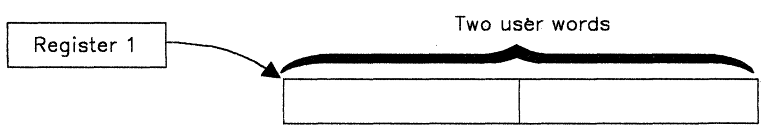

Figure 6-2. User Sampler Input Parameter Structure

Note: The user sampler must reside in SYS1.LPALIB. See "Adding Your Routines to RMF" later in this chapter.

Interval Processing

The interval processing user exit is ERBMFDUC. It is invoked at the start of the Monitor I session and at the end of each RMF interval.

The exit is invoked by IGX00022. When the exit gets control, register 1 points to a two-word address list. The first address points to the two-word area reserved for use by your routines. When these routines include a user sampler, one of these words, selected by your installation, will contain the address of the sampled data. The second address points to the RMF phase parameter. This parameter is a full word that contains X'4' when the exit is called during Monitor I session initialization, X'8' when the exit is called at the end of an RMF interval, or X'C' when the exit is called at the end of an RMF interval for which data collection was skipped. Figure 6-3 illustrates the input parameter structure.

When the interval processing exit routine is entered, the system is in supervisor state, and all interrupts are enabled. The routine runs in key 0. The routine can process the data generated by the user sampler. It can also collect its own data from system control blocks or system state indicators and format an SMF record. The SMF record can be written to the SMF data set; see *System Management Facilities (SMF)* for details on using the SMFEWTM macro instruction to write a user SMF record.

The SMF record or a record your routine formats as agreed by convention between ERBMFDUC and ERBMFRUR (the report writer exit routine) can be printed by your report writer. Your routine can format SMF record output, report record output, or both. When your routine formats any records to be printed by your report writer, the address of the formatted records must be placed in the user word selected by your installation. Because the user words are passed to your report writer, the records can then be printed in a formatted report.

When the length of the RMF interval exceeds 99 minutes, which can occur when RMF is not dispatched at the end of an interval, data collection for the interval is skipped. Because there is no data collected, RMF does not call the report writer user exit (ERBMFRUR); instead, ERBMFDUC is called twice. The phase parameter is X'8' for the first invocation of the exit routine and X'C' for the second. When the exit routine is called with a phase parameter of X'C', your routine must free the storage areas normally freed by ERBMFRUR. RMF issues a message to notify the operator that data collection was skipped for the interval.

When your routine has completed processing, return control by branching on register 14.

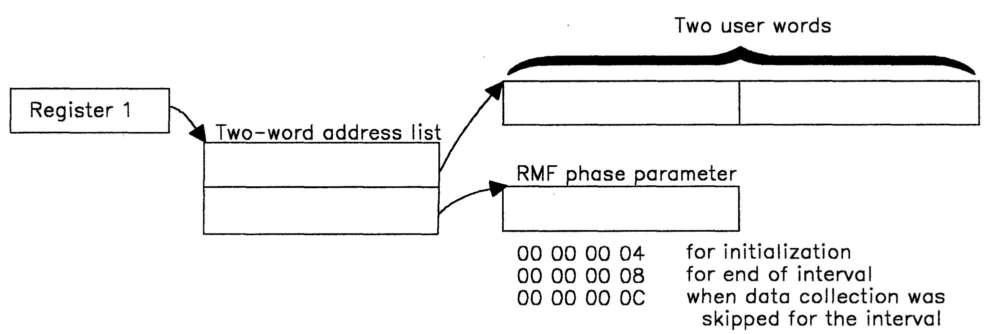

Figure 6-3. ERBMFDUC Input Parameter Structure

Report Writing During Session Processing

The report writer exit is ERBMFRUR. It is called once during the Monitor I session report writing phase.

The exit is invoked by ERBMFRGM during the report writing phase of interval processing. When the exit gets control, register 1 points to a two-word address list. The first address points to the two-word area reserved for use by user routines. The second address points to the RMF phase parameter, which is always X'10' for the report writer. Figure 6-4 illustrates the input parameter structure.

When the report writer exit is entered, the system is in problem state, and all interrupts are enabled. The routine runs in the user key 8. The user word your installation selects contains the address of the formatted records built by ERBMFDUC. Because all of your installation's exit routines use these words, the report writer must not alter their contents. Report writer processing must obtain output space for the printed reports, then write the reports for subsequent printing.

Before terminating, the routine must free the storage that contained the records formatted by ERBMFDUC.

When the report writer completes its processing, return control by branching on register 14.

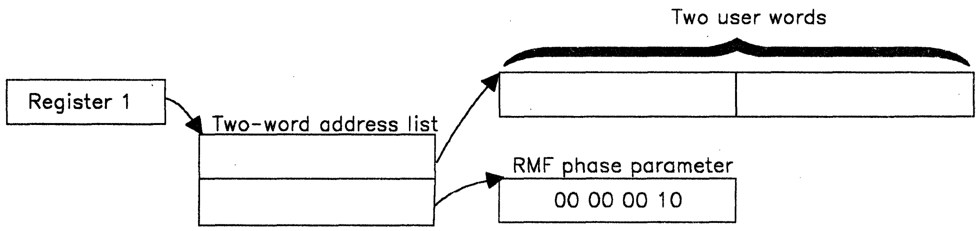

Figure 6-4. ERBMFRUR Input Parameter Structure

Termination

The termination exit is ERBMFTUR. It is called when the Monitor I session is terminated.

The exit is invoked by ERBMFTMA. When the exit gets control, register 1 points to a two-word address list. The first address points to a two-word area reserved for use by your routines. The second address points to the RMF phase parameter, which is always X'C' for termination. Figure 6-5 illustrates the input parameter structure.

When the termination routine is entered, the system is in supervisor state, and all interrupts are enabled. The routine runs in key 0. You would use this exit to page free any user samplers or data areas and to free any user SQA data areas obtained by the other exits, with one exception: during termination processing, ERBMFTUR gets control before the report writer exit (ERBMFRUR). Therefore, it must free only the SQA and global storage the other user routines obtained, but it **must** not free the storage the interval processing routine (ERBMFDUC) used to build records to be passed to the report writer. The address of this storage will be in the user word selected by your installation.

When the termination routine has completed processing, return control by branching on register 14.

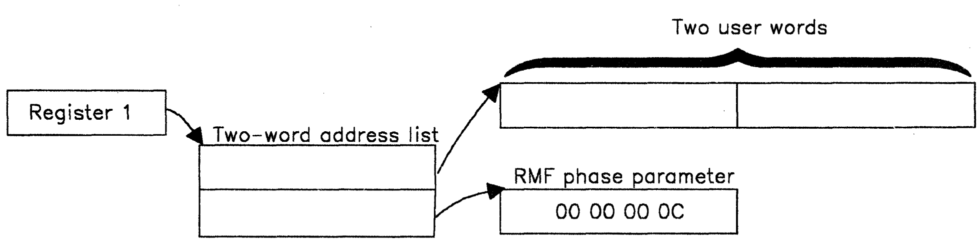

Figure 6-5. ERBMFTUR Input Parameter Structure

Tracing Your Own Field

Whenever the TRACE option is in effect during a Monitor I session, you can use the tracing routines to trace the contents of any SQA, CSA, or nucleus field that you require. The EXITS option, required to generate your own reports, is not required to use the trace facilities because the trace exit, ERBTRACE, is called whenever the TRACE option is in effect.

The field that you want to trace must be fixed in CSA, SQA, or the nucleus; it cannot contain negative values; and it must be from one to four bytes in length. Once you have selected your field, there are two steps required to enable RMF to trace the contents of the field. After you have performed these steps, you can then specify the name in the field name portion of the TRACE option. The steps can be performed in any order, but both must be done before you can use RMF to trace the field.

Step 1 -- Defining the Name to RMF

To define the name, you must add four fields to the RMF CSECT ERBMFTTB, which contains the names RMF recognizes as valid for tracing. The fields you must add are:

1. The name of the field to be traced. The name can be from one to eight bytes long. It must not be the same as any name already recognized by RMF. When the name of the field is less than eight bytes long, it must be padded on the right with blanks to a length of eight bytes.
2. The length of the name. This field is one byte long; the value must be from 01 to 08.
3. A one-byte constant that always contains the value X'DC'.
4. The length of the field to be traced. This field is one byte long; the value must be from 01 to 04.

One byte of binary zeroes must follow the last entry to be added; the byte of binary zeroes indicates the end of the variable-length trace table. Figure 6-6 shows an example of how to superzap ERBMFTTB to add a new name for tracing. The parenthesized numbers in the text refer to the parenthesized numbers in the figure. The example adds a nucleus field named MYDATA (1) that is two bytes in length (4) to the list of names valid for RMF tracing. The name is six bytes long (2), and the required constant is also supplied (3). A byte of binary zeroes (5) indicates the end of the trace table.

Adding the name definition to ERBMFTTB causes RMF to pass the name to ERBTRACE during each tracing cycle. The four fields must be added for each name you want RMF to trace; only the last entry must be followed by the byte of binary zeroes.

```
//ZAP       JOB         MSGLEVEL=1
//SS        EXEC        PGM=AMASPZAP
//SYSPRINT  DD          SYSOUT=A
//SYSLIB    DD          DSN=SYS1.LINKLIB,DISP=SHR
//SYSIN     DD          *
    NAME    ERBMFMFC        ERBMFTTB
    VER     040C            0040D7C1
    REP     040C            D4E8C4C1E3C14040 (1)
    REP     0414        06  (2)
    REP     0415        DC  (3)
    REP     0416        02  (4)
    REP     0417        00  (5)
/*
```

Figure 6-6. Example of Adding Name a to ERBMFTTB

Step 2 -- Replacing ERBTRACE

The tracing user exit is ERBTRACE. The function of ERBTRACE is to return to RMF the address of a valid user field. It is called by the RMF tracing routine whenever it encounters a trace name that is not the name of a field in the SRM domain table. To trace your own field, you must replace ERBTRACE with your own routine and link edit your ERBTRACE with the RMF CSECT ERBMFITR.

When ERBTRACE gets control, register 1 points to a two-word address list. The first address points to an eight-byte field that holds the name to be validated. The second address points to a full word to be used by ERBTRACE to return the address of the user field to RMF. Figure 6-7 illustrates the input parameter structure.

When ERBTRACE is entered, the system is in supervisor state, and all interrupts are enabled. The routine runs in key 0. It must examine the field name passed to it by RMF to determine if the name is a user field name. When the name is a valid user name, place the address of the field to be traced in the first parameter, set a return code of zero in register 15, and return control. If the name is not one recognized as a valid user name, always set a non-zero return code in register 15 before returning control. The non-zero return code tells RMF to process the name.

When your processing is completed, return control by branching on register 14.

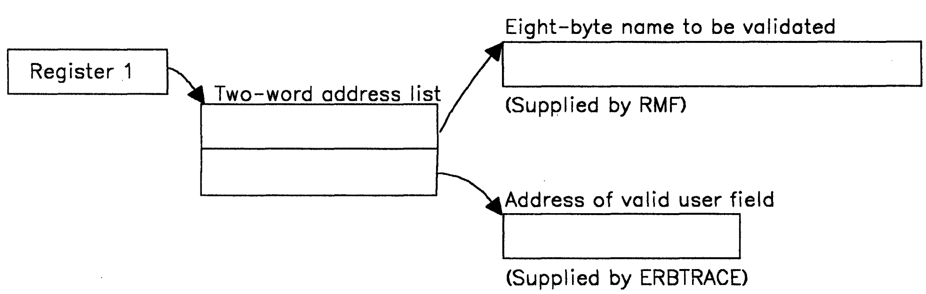

Figure 6-7. ERBTRACE Input Parameter Structure

Report Writing by the Post Processor

The post processor user exit is ERBMFPUS. It is called during post processing at the point when the post processor first encounters each SMF record.

The exit is invoked by ERBRMFPP, the post processor control routine, whenever the EXITS control statement is specified. When the exit gets control, register 1 points to a two-word address list. The first address points to a full word that contains the address of the SMF record to be processed. The second address points to a full word reserved for the use of your routine. The user word contains zeros the first time the exit is called, and the post processor does not modify its contents between invocations of the user exit routine. Thus, the word can be used to save information, such as the address of a DCB, that is needed by a subsequent invocation of the user exit routine. Figure 6-8 illustrates the input parameter structure.

When the post processor user exit is entered, the system is in problem state and all interrupts are enabled. The routine runs in the user key 8.

Your routine examines the SMF record passed to you, performs any required processing, and set a return code in register 15. The return code depends on the action you want the post processor to take. A return code of 0 tells the post processor to continue processing the SMF record. A return code of 4 tells the post processor to ignore the SMF record; set a return code of 4 when the exit routine has, for example, processed the record or determined that it should not be processed. A return code of 8 indicates that the post processor should terminate.

The processing your exit performs can consist of formatting the data in the records the interval processing user exit routine (ERBMFDUC) generates into a printed report. Your exit could also screen the SMF records that the post processor encounters to determine which records are to be included in any reports generated by the post processor, or it could use the SMF records RMF generates as input to your own report. Because all SMF records are passed to the user exit, ERBMFPUS could also be used to incorporate any SMF data reduction routines used at your installation into the RMF post processing function.

When your routine has finished processing, set the appropriate return code in register 15 and return control to the RMF post processor by branching on register 14.

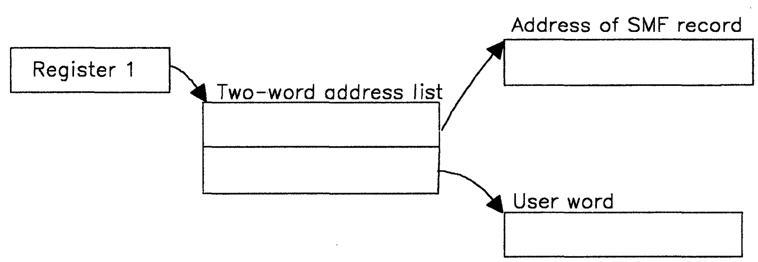

Figure 6-8. ERBMFPUS Input Parameter Structure

Adding Your Routines to RMF

Before your Monitor I session user exit routines can be tested and used, they must be assembled and link edited with the appropriate RMF modules. Figure 6-9 shows sample JCL for performing the required link edit for all user routines except the sampler routine. If you have a user sampler, a separate link edit is required; a sample is shown in Figure 6-10.

```
//LINKEXIT      JOB   MSGLEVEL=1
//STP           EXEC  PGM=IEWL,PARM='link edit parameters'
//SYSPRINT      DD    SYSOUT=A
//SYSLMOD       DD    DSN=SYS1.LINKLIB,DISP=(OLD,KEEP)
//SYSUT1        DD    UNIT=SYSDA,DISP=(,DELETE),
//                    SPACE=(TRK,(20,5))
//SYSLIN        DD    *
    (ERBMFIUC object deck)
    ENTRY   ERBMFIUC
    NAME    ERBMFIUC(R)
    (ERBMFDUC object deck)
    ENTRY   ERBMFDUC
    NAME    ERBMFDUC(R)
    (ERBMFRUR object deck)
    ENTRY   ERBMFRUR
    NAME    ERBMFRUR(R)
    (ERBMFTUR object deck)
    ENTRY   ERBMFTUR
    NAME    ERBMFTUR(R)
    (ERBTRACE object deck)
    INCLUDE SYSLMOD(ERBMFITR)
    ENTRY   ERBMFITR
    NAME    ERBMFITR(R)
    (ERBMFPUS object deck)
    ENTRY   ERBMFPUS
    NAME    ERBMFPUS(R)
/*
```

Figure 6-9. Replacing User Exits

```
//LINKEXIT      JOB   MSGLEVEL=1
//STP           EXEC  PGM=IEWL,PARM='link edit parameters'
//SYSPRINT      DD    SYSOUT=A
//SYSLMOD       DD    DSN=SYS1.LPALIB,DISP=(OLD,KEEP)
//SYSUT1        DD    UNIT=SYSDA,DISP=(,DELETE),
//                    SPACE=(TRK,(20,5))
//SYSLIN        DD    *
    (user sampler object deck)
    ENTRY   entry name
    NAME    sampler name
/*
```

Figure 6-10. Adding a User Sampler

Monitor II Session User Reports

RMF generates a Monitor II session report by invoking a data-gathering module and a data-reporting module in response to either:

- a menu item identifying a display session report
- an option identifying a background session report

From an external viewpoint, the menu item and the option are different because they are used during different types of sessions, have slightly different syntax, and produce either display output or printed output. However, from an internal point of view, the menu item and the option are very similar. The valid menu items for a display session are listed in the RMF CSECT ERBFMENU. The valid options for a background session are listed in the RMF CSECT ERBBMENU. The formats of the entries in each list are identical. When an option or menu item is specified during a session, RMF uses the data entry for the report in the list appropriate for the session type to verify that the option or menu item is valid and to load the required data gatherer and data reporter modules.

Each list contains an entry called USER that enables you to add a single user report. When USER is specified, RMF loads modules ERBGUS99, the data gatherer for USER, and ERBRUS99, the data reporter for USER. By replacing these two modules with your own routines, you can add a single report to the Monitor II reports provided by RMF. This process is described later in this chapter under "Coding a User Report."

The data gathering module and the data reporting module communicate through a type 79 SMF record. The data gatherer formats the record and completes the required data fields. The data reporter uses the data in the record to generate a formatted report for printing or display. See "SMF Record Type 79."

To add more then one Monitor II session report, you must, in addition to providing a data gatherer and a data reporter, add an entry to ERBFMENU for a display session report and to ERBBMENU for a background session report. Then, when your option or menu item is specified during a session, RMF will load your data gatherer and data reporter to generate the report. The process to follow to add an entry to the option list and menu list is described later in this chapter under "Installing a User Report."

Guidelines

Each of the user functions is described in detail in the following sections. The following guidelines apply to all Monitor II user exit routines.

- All of the user exit routines must be reenterable.
- All user-written exit routines receive control in 31-bit addressing mode.
- The routines must save registers when they receive control and restore registers when they return control. Register 13 contains the address of the register save area; register 14 contains the return address; and register 15 contains the entry address.
- All of the user exit routines receive control in problem state, key 8.

SMF Record Type 79

SMF record type 79 must be used to record data gathered by a user data gathering routine. Figure 6-11 shows the layout of the record sections that are common to all Monitor II data gatherers, whether coded by a user or provided by RMF. The figure illustrates the layout of these common sections by showing the expansion of the RMF mapping macro ERBSMF79.

The fields in the common sections fall into three categories. Each category is indicated by a letter in the figure that corresponds to the letters in the following text:

(A) The fields that the RMF routines fill in before the data gathering routine is invoked.

(B) The fields that the data gathering routine must fill in during its processing. (See "Relocate Blocks" later in this section.)

(C) The fields that the RMF routines will fill in when the RECORD option is in effect. RMF completes these fields after the data gatherer returns control but before the record is written to the SMF data set. During a display session or a background session when NORECORD is in effect, these fields are not completed because the record is not actually written to the SMF data set.

Before invoking the data gatherer, RMF calculates the length of the storage buffer required for the record, as described later under "Relocate Blocks," obtains a buffer for the record, and fills in some of the common section fields. The address of the SMF record buffer is passed to the data gatherer. The data gatherer fills in some fields in the common section and all of the data section of the record.


```
******************** COMMON SMF HEADER *****************************

    SMF79HDR   DSECT
(C) SMF79LEN    DS      BL2     RECORD LENGTH
    SMF79SEG   DS      BL2     SEGMENT DESCRIPTOR
(C) SMF79FLG    DS      BL1     HEADER FLAG BYTE
    SMF79RRF   EQU     X'80'   NEW SMF RECORD FORMAT IF=1
    SMF79SUT   EQU     X'40'   SUBTYPE UTILIZED IF=1
    SMF79VXA   EQU     X'04'   MVS/XA IF=1
    SMF79OS    EQU     X'02'   OPERATING SYSTEM IS OS/VS2
(C) SMF79RTY    DS      BL1     RECORD TYPE
    SMF79TME   DS      BL4     TOD RECORD WRITTEN
    SMF79DTE   DS      PL4     DATE RECORD WRITTEN
(C) SMF79SID    DS      CL4     SYSTEM ID FROM INSTALLATION
(C) SMF79SSI    DS      CL4     SUBSYSTEM ID (RMF)
(B) SMF79STY    DS      BL2     SUBTYPE
(A) SMF79TRN    DS      BL2     NUMBER OF TRIPLETS IN THIS RECORD
               DS      BL2     RESERVED
(A) SMF79PRS    DS      BL4     OFFSET TO RMF PRODUCT SECTION
(A) SMF79PRL    DS      BL2     LENGTH OF RMF PRODUCT SECTION
(A) SMF79PRN    DS      BL2     NUMBER OF RMF PRODUCT SECTIONS

  ********************** INDIVIDUAL HEADER EXTENSION *******************

(A) SMF79MCS    DS      F -     OFFSET TO MONITOR II CONTROL SECTION
(A) SMF79MCL    DS      H -     LENGTH OF MONITOR II CONTROL SECTION
(A) SMF79MCN    DS      H -     NUMBER OF MONITOR II CONTROL SECTION
(B) SMF79ASS    DS      F -     OFFSET TO DATA SECTION
(B) SMF79ASL    DS      H -     LENGTH OF DATA SECTION
(B)(A) SMF79ASN  DS      H -     NUMBER OF DATA SECTION
(A) SMF79DCS    DS      F -     OFFSET TO DATA CONTROL SECTION
(B) SMF79DCL    DS      H -     LENGTH OF DATA CONTROL SECTION
(B) SMF79DCN    DS      H -     NUMBER OF DATA CONTROL SECTION
    SMF79QSS   DS      F -     OFFSET IOQ GLOBAL SECTION
    SMF79QSL   DS      H -     LENGTH IOQ GLOBAL SECTION
    SMF79QSN   DS      H -     NUMBER IOQ GLOBAL SECTION
```

Figure 6-11 (Part 1 of 3). ERBSMF79 Mapping Macro Expansion

```
    *********************  COMMON SMF PRODUCT SECTION *******************

     SMF79PRO   DSECT
 (C) SMF79MFV    DS       CL2      RMF VERSION # IN EBCDIC
 (C) SMF79PRD    DS       CL8      PRODUCT NAME
     SMF79IST   DS       PL4      TOD MONITOR 1 INTERVAL START: 0HHMMSSF
 (C) SMF79DAT    DS       PL4      DATE MONITOR 1 INTERVAL START: 00YYDDDf
     SMF79INT   DS       PL4      DURATION OF MONITOR 1 INTERVAL: MMSSTTTF
     *
     SMF79MFL   DS       XL2      RECORD MAINTENANCE INDICATION.
    *                            THIS FIELD HAS A DIFFERENT MEANING FOR...
    *                            ...DIFFERENT SMF RECORD TYPES
    *
 (B) SMF79SAM    DS       BL4      NUMBER OF SAMPLES
     SMF79RV2   DS       BL2      RESV
     SMF79FLA   DS       BL2      FLAGS
     SMF79CNV   DS       X'80'    DATA CONVERT.FROM VER.2
     SMF79ISS   DS       X'40'    INVALID SAMPLES TO BE SKIPPED
 (C) SMF79RLS    DS       CL4      OS/VS2 RELEASE NUMBER IN EBCDIC
 (B) SMF79CYC    DS       PL4      CYCLE IN PACKED DECIMAL 000TTTTF
 (B) SMF79MVS    DS       PL8      MVS SOFTWARE LEVEL
 (B) SMF79IML    DS       PL1      RMF I/O MEASUREMENT LEVEL
 (B) SMF79PRF    DS       XL1      PROCESSOR FLAGS
 (B) SMF79QES    EQU      X'80'    EQUIPPED WITH EXTENDED STORAGE

    *********************  MONITOR II CONTROL SECTION  *******************

     R79CHL     DSECT             COMMON RECORD 79 HEADER
 (B) R79GTOD     DS       XL4 -    DATA GATHERER CALL TOD
 (B) R79LF2      DS       XL1 -    FLAG BYTE
     R79PAR     EQU      X'80'    NOT ENOUGH RELOCATE SECTION TO
     *                            COMPLETE DATA GATHERING
     R79SG      EQU      X'40'    REPORT TO BE SORTED BY SG
     R79RV1     DS       XL1 -    RESERVED
 (C) R79SES      DS       CL2 -    SESSION NAME
     R79RSV     DS       XL2 -    RESERVED
     R79USER    DS       XL2 -    USER FIELD
 (C) R79RID      DS       CL8 -    MEASUREMENT NAME
 (C) R79CTXTL    DS       XL2 -    LEN OF COMMAND TEXT
 (C) R79CTEXT    DS       CL32 -   COMMAND TEXT
 (C) R79DTXTL    DS       XL2 -    LEN OF DEFAULT DR TEXT
 (C) R79DTEXT    DS       CL32 -   DEFAULT DR TEXT
 (C) R79IST      DS       CL4 -    MON III INTERVAL START TIME :0HHMMSSF
```

Figure 6-11 (Part 2 of 3). ERBSMF79 Mapping Macro Expansion


```
********************** DATA  SECTION  *********************************

R799LCU    DS      BL2        LOGICAL CONTROL UNIT NUMBER 0 TO 255
R799SGN    DS      CL8        STORAGE GROUP NAME
```

Figure 6-11 (Part 3 of 3). ERBSMF79 Mapping Macro Expansion

Relocate Blocks

The data section of SMF record type 79 is unique to each report. It is composed of one or more data sections called **relocate** blocks and, possibly, one data control section. A relocate block is the portion of the SMF record that contains the data for one report data line. A record for a row report has one relocate block. A record for a table report has multiple relocate blocks; for example, the SMF record for the address space state data report includes one relocate block for each address space included in the report. When your SMF record has multiple relocate blocks and you are gathering data that applies to all of them, you can, instead of reporting the data in each relocate block, place this common data in a data control section, as described later under "Data Control Section."

The format of the data in the relocate block depends on the report you are generating. You set the format that best meets your needs. When you are generating a table report, the SMF record consists of multiple relocate blocks, and each relocate block must have the same length.

When you add a menu item to ERBFMENU or an option to ERBBMENU, the entry that describes the new report must include a field that specifies the length of the relocate block, the maximum number of possible relocate blocks, and the length of the data control section. For information on how to add an entry to ERBFMENU or ERBBMENU, see "Using the PICTURE Macro" later in this chapter. To determine the storage to allocate, RMF multiplies the length of the relocate block by the maximum number of relocate blocks and adds this value to the length of the data control section and the common section. The result of this computation is the maximum possible length of the SMF record, and RMF allocates a buffer for the record that is equal in size to the maximum length.

To determine the actual length of the SMF record, the data gatherer must complete the fields in the individual header extension section that describe the offset, length, and number of data sections and the data control sections. After the data gatherer has completed its processing and returned control, RMF uses these values to determine the length of the SMF record to be written to the SMF data set, a calculation that is performed only when the RECORD option is in effect for a background session. Note that the value your routine sets in SMF79ASL and the value specified for RBLEN in the PICTURE macro for the report should be identical.

Other fields in the common section that the data gather completes are R79GTOD and SMF79STY. R79GTOD must contain a packed decimal value that indicates the time when the data gatherer was invoked, in the form 0hhmmssF, where F is the sign. SMF79STY can contain the subtype number of the SMF record that you are creating. You use this number as a unique identifier for each record subtype that you create; no subtype number should be less than 1000.

The maximum length of an SMF record is 32,756 bytes; any records that exceed this length are truncated before they are written to the SMF data set. Truncation, which can occur only during a background session when the RECORD option is in effect, occurs at the last relocate block boundary within the maximum length. When truncation occurs, RMF adjusts the field indicating the capacity of the buffer (SMF79ASN) to indicate the actual number of relocate blocks in the record. If no truncation occurs, RMF leaves SMF79ASN unchanged.

Data Control Section

A data control section is useful when your SMF record might have many relocate blocks and some of the data you are gathering is common to all of them. For example, the channel path Monitor II control section (subtype C) uses a control section to record the number of times the channel was sampled. To use a data control section:

1. Set the value for the FBLEN parameters on the PICTURE macro instruction for your report, as described later under "Using the PICTURE Macro."
2. Format the data control section to hold the common data.
3. Place it between the Monitor II control section and the data section. SMF79DCS contains the offset at which it should start.
4. Set SMF79DCL and SMF79DCN to the length and number of the data control sections.
5. Set the offset to the first data section SMF79ASS to point to the end of the data control section.

When a data control section is *not* used:

1. Set SMF79ASS to the value in SMF79DCS.
2. Set SMF79DCL and SMF79DCN to 0.

Coding a User Report

To add a Monitor II report, you must code your own data gatherer module and data reporter module. These modules can reside in SYS1.LINKLIB, SYS1.LPALIB, a steplib, a joblib, a tasklib, or a library in a linklist.

The primary means of communicating data between the gatherer and the reporter is the type 79 SMF record. The gatherer collects data from whatever areas it can access (it runs in problem state with a key of 8) and places the data in the SMF record. The reporter takes the data from the SMF record, formats it for output, and passes it to the RMF putline routine. During a Monitor II background session, the data reporter would be called when the REPORT option is in effect. When NOREPORT and RECORD are in effect, RMF writes out the SMF records that the data gatherer formats, and the data reporter is not invoked. Your data reporter can be invoked at a later time by the post processor.

A Monitor II session report can have operands that the report user specifies when he requests the report. Any operands specified when a report is requested are passed to both the data gatherer and the data reporter. The defaults established for each possible operand are specified in the option list or menu list entry for the report; these defaults are also passed to both the data gatherer and the data reporter. Your routines can also include hard-coded default operands.

Because the option list and menu item list are in different RMF control sections, you can set different default operands for a background session and a display session. Each list entry contains separate fields for the data gatherer default operands and the data reporter default operands; you can thus set different default operands for the data gatherer and the data reporter. For example, the default operands for the RMF address space state data gatherer module cause data to be gathered on all address spaces in the system; to limit the actual output produced, the defaults for the reporter cause only the active address spaces to be reported. "Using the PICTURE Macro" describes how to specify default operands.

RMF passes parameters to both the gatherer and reporter; these parameters include a subpool number that indicates the subpool from which the routines should obtain the storage they require, and two user words that can be used for communication between the data gatherer and the data reporter. Because the same two words are passed to both routines, use of these words must be governed by conventions established by your installation.

Note: A system status line precedes each display report supplied by IBM. RMF obtains the data for this line before it invokes the data gatherer for the report. RMF will generate the same system status line before each user-coded display report.

Data Gatherer

The data gatherer runs in problem state, with a key of 8, and in 31-bit addressing mode. The data gatherer must be reenterable. It receives control by a BALR instruction and must save the registers when it receives control and restore the registers when it returns control. Register 13 contains the address of the register save area; register 14 contains the return address; and register 15 contains the entry address.

Upon entry to the data gatherer, register 1 points to a contiguous list of seven addresses that point to seven input parameters. The first address points to the first parameter, the second address points to the second parameter, and so forth. The input parameters are:

First Parameter: A fullword entry code that must always be X'2'.

Second Parameter: The operands, if any, specified by the report user when he requested the report, in the form:

LL	text

LL
A two-byte length field indicating the length of the following text (does not include the two bytes of LL)

text
A character string of up to 32 characters containing the input operands

When the report has no operands or the report request did not include operands, LL is set to zeros.

RMF determines the operands to be placed in **text** by scanning the report request. The first non-blank character after the report name is assumed to be the first character of the operand field. The next blank character is assumed to mark the end of the operand field.

Third Parameter: The default operands from ERBFMENU or ERBBMENU, in the form:

LL	text

LL
A two-byte length field indicating the length of the following text (does not include the two bytes of LL)

text
A character string of up to 32 characters containing the default operands

When the report has no operands or no default operands, LL is set to zeros.

Fourth Parameter: The pointer to the SMF record buffer where your routine is to place the data it gathers.

Fifth Parameter: The first of the two words reserved for the use of your routines.

Sixth Parameter: The second of the two words reserved for the use of your routines.

Seventh Parameter: A byte containing the number of the subpool to use when you issue a GETMAIN to obtain the storage your routine requires.

The processing your data gathering routine performs is determined largely by the nature of the report for which you are gathering data. This processing should include a validation of the entry code in the first parameter to verify that it is X'2'. If it is not, set a return code of 8 in register 15 and return control.

If the report has operands that can be specified when the report is requested, check the second input parameter to determine if the request specified operands. If it did, validate the syntax of the operands; if the syntax is invalid, set a return code of 4 in register 15 and return control. If the request did not specify operands, verify the syntax of the default menu operands passed as the third input parameter; if the syntax is invalid, set a return code of 24 in register 15 and return control.

Your routine should complete the required fields in the SMF record common section (the **(B)** fields in Figure 6-11), using the RMF mapping macro ERBSMF79 to access the fields in the common section. The address of the storage buffer obtained for your record is passed in the fourth input parameter. Your routine would gather the data required and format the data section of the record as agreed upon by convention between the data gatherer and the data reporter. Should your routine locate no data that is applicable to the report requested, set a return code of 16 in register 15 and return control.

When your routine has finished processing, set a return code in register 15 and return to the caller by branching on the contents of register 14. Figure 6-12 shows the possible return codes, their meaning, and the action RMF takes in response. These return codes apply to both the data gatherer and the data reporter.

Note: If your report will be run only during a display session, you can perform both the data gathering function and the data reporting function in the data reporter module. In this case, your data gatherer's only function would be to set a return code of zero in register 15. However, if you choose to perform both functions in the data reporter module, your report cannot run during a Monitor II background session and, during a display session, you will not be able to use the recall command to re-display your report.

Code	Meaning	RMF Response (Display Session)	RMF Response (Background Session)
0	Successful completion	The session continues.	The session continues.
4	Invalid operand syntax	The command is displayed as entered.	Message ERB409I is issued. The current measurement continues if the error was detected by the data reporter and RECORD is in effect; otherwise, the measurement is discontinued. The session continues. The operator can modify the session options.
8	Invalid entry code	Abend. The user code is 1402	Abend. The user code is 1402.
12	I/O error	Messages ERB403I and ERB404I are displayed, including the SYNAD text.	The current measurement continues when RECORD is in effect, but no subsequent reports are printed; otherwise, the measurement is discontinued. The session continues.
16	No data found	Message ERB405I is displayed.	Message ERB405I is issued. No report or SMF record is produced for this interval. All measurements continue.
20	ESTAE macro failed	Message ERB406I is displayed.	Message ERB406I is issued. The current measurement continues if the error was detected by the data reporter and RECORD is in effect; otherwise, the measurement is discontinued. The session continues.
24	Menu default operand syntax error	Message ERB407I is displayed, including the menu defaults and advice to retry the report, specifying all operands.	Message ERB407I is issued. The current measurement continues if the error was detected by the data reporter and RECORD is in effect; otherwise, the measurement is discontinued. The session continues.
28	The amount of data to be gathered exceeds the number of available relocate blocks.	Message ERB411I is displayed.	Message ERB411I is issued. The report or SMF record produced for the interval includes only the data gathered before the condition was detected. All measurements continue.
32	Monitor I report not active	Message ERB412I is displayed.	Message ERB412I is issued. No report or SMF record is produced for the interval. All measurements continue.
36	Monitor I interval is less than Monitor II interval	Message ERB413I is displayed.	Message ERB413I is issued. No report or SMF record is produced or the interval. All measurements continue.
40	The SRM's store channel path status facility is not active. Used by channel path activity (CHANNEL) report	Message ERB264I is displayed.	Message ERB264I is issued. No report or SMF record for channel path activity is produced; the current measurement is discontinued. All other measurements continue.
>41	Unexpected	Message ERB408I is displayed.	Message ERB408I is issued. The current measurement continues if the error was detected by the data reporter and RECORD is in effect; otherwise, the measurement is discontinued. The session continues.

Figure 6-12. Return Codes from the Data Gatherer and Data Reporter

Data Reporter

The data reporter runs in problem state, with a key of 8, and in 31-bit addressing mode. The data reporter must be reenterable. It receives control by a BALR instruction and must save the registers when it receives control and restore the registers when it returns control. Register 13 contains the address of the register save area; register 14 contains the return address; and register 15 contains the entry address.

The data reporter formats each line in the report, using the data placed in the type 79 SMF record by the data gatherer. The RMF putline routine is used to perform the actual output operation.

Because the putline routine handles the actual output operations, your data reporter can function identically during a background session, a display session, a display session in hardcopy mode, or an execution of the post processor. The putline routine writes the line to a logical screen buffer for a display session, to a logical screen buffer and an output data set for a display session in hardcopy mode, or to an output data set for a background session or an execution of the past processor. For a display session, the screen is updated to show the lines collected by the putline routine when your data reporter returns control. Note that RMF handles any framing required for the display session user to view all the frames in a multi-frame table report after the data reporter completes its processing.

The data reporter you code can generate either a row report or a table report. The maximum number of header lines is two.

A row report consists of one or two header lines and a single data line. For a row report, RMF invokes the data reporter twice: once to format the header line(s) and once to format the data line. When a row report is executed repetitively, RMF invokes the reporter to format the header line(s) for the first execution; for all subsequent executions, the reporter is invoked to format a data line.

A table report consists of one or two header lines and a variable number of data lines. For a table report, RMF invokes the data reporter once to format both the header line(s) and the data lines. The number of data lines must be less than or equal to the number of relocate blocks created in the SMF record by the data gatherer.

Upon entry to the data reporter, register 1 points to a contiguous list of eleven addresses that point to eleven input parameters. The first address points to the first parameter, the second address points to the second parameter, and so forth. The input parameters are:

First Parameter: A full word entry code that can be either X'1' or X'2'. X'1' indicates that the reporter is to format the header line(s) for a row report. X'2' indicates, for a row report, that the reporter is to format the single data line. For a table report, the entry code should always be X'2', indicating that the reporter is to format both the header line(s) and the data lines.

Second Parameter: A full word report mode indicator that can have either of the following values:

X'1'
Total mode; the values in the report are to reflect session totals.

X'2'
Delta mode; the values in the report are to reflect changes since the last request for the report.

Third Parameter: The operands, if any, specified by the report user when he requested the report, in the form

LL	text

LL
A two byte length field indicating the length of the following text (does not include the two bytes of LL).

text
A character string of up to 32 characters containing the report operands.

When the report has no operands or the report request did not include operands, LL is set to zeros.

Fourth Parameter: The default operands from ERBFMENU or ERBBMENU, in the form:

LL	text

LL
A two byte length field indicating the length of the following text (does not include the two bytes of LL).

text
A character string of up to 32 characters containing the default operands.

When the report has no operands or no default operands, LL is set to zeros.

Fifth Parameter: The address of the current SMF record buffer; that is, the buffer where the data gatherer has placed the data for the current execution of the reporter.

Sixth Parameter: The address of the previous SMF record buffer; that is, the buffer where the data gatherer placed the data for the previous execution of the report. When the report mode (the second parameter) indicates delta mode, the data fields in the previous SMF record enable your data reporter to calculate the changes that have occurred since the last request for the report.

Seventh Parameter: The first of the two words reserved for the use of your routines.

Eighth Parameter: The second of the two words reserved for the use of your routines.

Ninth Parameter: A byte containing the number of the subpool to use when you issue a GETMAIN to obtain the storage your routine requires.

Tenth Parameter: The address of the RMF putline routine. When the data reporter has formatted a report line, it calls the putline routine to perform the actual output operation.

Eleventh Parameter: The control block address that your data reporter must pass to the putline routine.

The processing your data reporting routine performs is determined largely by the nature of the report for which you are formatting report lines. This processing should include a validation of the entry code. If it is not a valid code, set a return code of 8 in register 15 and return control. If your report is a row report, examining the entry code determines whether your routine has been invoked to format the header line(s) or the data line for the report.

If the report has operands that can be specified when the report is requested, check the third input parameter to determine if the request specified operands. If it did, validate the syntax of the operands; if the syntax is invalid, set a return code of 4 in register 15 and return control. If the request did not specify operands, verify the syntax of the menu default operands passed as the fourth input parameter; if the syntax is invalid, set a return code of 24 in register 15 and return control.

If your report contains fields that are affected by the session mode – either delta mode or total mode – check the second input parameter to determine which mode is in effect. When delta mode is in effect, use the data fields in the previous SMF record buffer (pointed to by the sixth parameter) and the data fields in the current SMF record buffer (pointed to by the fifth parameter) to calculate the changes that have occurred since the last report request.

When your routine has formatted a report line, it should invoke the RMF putline routine to perform the actual output operation. To use the putline routine, perform the following steps:

1. Set up the input parameters that the putline routine requires. To do this, set register 1 to point to a list of four addresses that point to the following four parameters:

   First Putline Parameter: The record you have formatted, preceded by a two-byte length field. The length specified **must not** include the two bytes of the length field. The maximum record length is 79 characters. Note that the 3270 field attribute bytes must **not** be included; RMF supplies these bytes.

   Second Putline Parameter: A two-byte field that tells the putline routine whether the record you have formatted is a header line or a data line. The field must contain one of the following:

   'HD'
      Indicates that the record is a header line

   'DT'
      Indicates that the record is a data line

   Header lines generally contain column headings. These lines are repeated when the terminal user frames forward through a multi-frame table report or when the hardcopy output crosses a page boundary.

Third Putline Parameter: A one-byte field; its bits have the following meaning:

Bit	Meaning
0	Set to 1 if high intensity display is desired. Set to 0 if low intensity display is desired. (The bit is ignored during a background session.)
1-7	Reserved. These bits must be set to zeros.

Fourth Putline Parameter: The control block address that RMF passed to your data reporter in the eleventh input parameter.

2. Invoke the putline routine using standard linkage conventions. Set register 13 to point to your register save area, set register 15 to the address of the putline routine (passed to your data reporter in the tenth parameter), and pass control to the putline routine by a BALR 14,15 instruction.
3. When the putline routine returns control to the data reporter, a return code is set in register 15. A return code of zero indicates successful completion. A return code of 4, indicates an uncorrectable I/O error; set a return code of 12 in register 15 and return control.

When your data reporter has finished processing, set a return code in register 15 and return control by branching on the contents of register 14. Figure 6-12 shows the possible return codes, their meaning, and the action RMF takes in response to each code.

Installing a User Report

Once your data gatherer and data reporter are coded, two steps are required to install the report:

1. Include an entry for the report in the option list for a background session (ERBBMENU) and the menu list for a display session (ERBFMENU), depending on the type of session during which your report can be run.

   If data collected during a Monitor II background session is to be reported during execution of the post processor, a copy of the option list control section (ERBBMENU) that includes the entry for your report must be link edited with the post processor.

   RMF supplies the PICTURE macro to simplify the process of adding or changing an entry in the option list or menu list. See "Using the PICTURE Macro" later in this chapter. You can also superzap an entry to make changes when the length of the entry is not changed.
2. Link edit your data gatherer and data reporter and test your report.

The option list or menu list consists of a set of variable-length entries, each describing a valid report. The option list appears in the RMF control section ERBBMENU; the menu list appears in ERBFMENU. Two separate control sections are provided to allow for a report that will run only during a background session or only during a display session. Also, the two different control sections allow different sets of default operands to be established for display sessions and background sessions. For example, you might want the display defaults to specify a limited set of possible data, while the background defaults specify all possible data.

The steps required to add an entry to the list are:

1. Determine whether the USER entry supplied by RMF is appropriate for your report. The USER entry contains specifications for a table report (RPTTYP=T) with a single relocate block (MAXRBS=1) that is four bytes long (RBLEN=4). The report title is 'USER PICTURE'. If the entry is not appropriate for your report, replace the entry with a new entry for USER.
2. If you are changing the USER entry or adding a new entry, make a copy of ERBFMENU for a display report or ERBBMENU for a background session — or both — from the source code data set.
3. In the copy you have made, either replace the USER entry or insert a new PICTURE macro. For a new display report, insert the PICTURE macro where you want the new report to appear in the menu frame. See "Using the PICTURE Macro" later in this section for details.
4. Assemble ERBFMENU for a display report and ERBBMENU for a background report.
5. Link edit the menu list or option list CSECT(s) that you have assembled into the RMF load modules. A sample of the control statements required is:

```
//SYSLMOD   DD SYS1.LINKLIB,DISP=(OLD,KEEP)
//SYSLIN    DD   *
     (ERBFMENU object deck)
     (ERBBMENU object deck)
     INCLUDE SYSLMOD(ERBMFMFC)
     ENTRY ERBMFMFC
     ALIAS ERBMFMPR
     SETCODE AC(1)
     NAME ERBMFMFC(R)
     (ERBFMENU object deck)
     INCLUDE SYSLMOD(RMFMON)
     ENTRY ERBMFTSO
     NAME RMFMON(R)
     (ERBBMENU object deck)
     INCLUDE SYSLMOD(ERBRMFPP)
     ENTRY ERBRMFPP
     NAME ERBRMFPP(R)
/*
```

Notes:

- If your report will run only during a TSO display session, the change to ERBMFMFC is not required.
- The change to ERBRMFPP is required only when your report will be printed during execution of the RMF post processor.

To install your report, you must link edit your data gatherer and data reporter.

If you are using the USER entry, name your gatherer routine ERBGUS99; name your reporter routine ERBRUS99. Replace the dummy RMF modules that have these names with your own routines. The link edit control statements required are:

```
(ERBGUS99 object deck)
ENTRY ERBGUS99
NAME ERBGUS99(R)
(ERBRUS99 object deck)
ENTRY ERBRUS99
NAME ERBRUS99(R)
```

If you are not using the USER entry, give your data gatherer and data reporter modules names that match the names you are specifying in the PICTURE macro for the report that you are adding. Link edit the modules as shown in the above control statements, replacing ERBGUS99 with the name of your data gatherer and ERBRUS99 with the name of your data reporter.

Once your modules have been link edited, you are ready to test your report. You might find it simpler to test your new report on TSO before making it available to other RMF users at your installation. Perform the following steps:

1. Use a testing tasklib, a special partitioned data set (for example, TESTLIB.LOAD). Place your data gatherer, data reporter, and the RMFMON load module that includes the new menu list in the testing tasklib.
2. You can then test the new report by entering:

   ```
   CALL TESTLIB(RMFMON)
   ```

   The new menu should appear on the screen in response to this command. You can then invoke your report by specifying its menu item name.

If your report routine terminates abnormally, you can obtain a dump by replying 'STOP' to the messages describing the abnormal termination.

Using the PICTURE Macro

The PICTURE macro describes a Monitor II session report to RMF. Use the PICTURE macro to replace the USER description or add or replace any entry in either ERBBMENU or ERBFMENU. The PICTURE macro is located in SYS1.RMFMAC01.

The syntax of the macro and the meaning of each operand are as follows:

```
[name]  PICTURE ID=name,
                GATHER=gathername,
                REPORT=reportname,
                RBLEN=length,
                RPTTYP={R|T}
                [,PFK=n]
                [,TITLE='title']
                [,DGTEXT='dgdefaults']
                [,DRTEXT='drdefaults']
                [,MAXRBS=nn]
                [,FBLEN=len]
```


ID = name
The option or menu item that will identify the report.

The name must consist of one to eight alphameric characters. The first character must not be 'R'; RMF takes 'R' to be a request to recall a report. For a display report, this name will appear on the menu frame.

GATHER = gathername
The name of the module RMF is to invoke to gather data for the report.

PFK = n
The PF key number associated with the report, where n is a one-digit or two-digit decimal identifier in the range of 1 to 24. For a display report, this number appears in the menu frame. If a PF key is not specified, the report is not associated with a PF key.

REPORT = reportname
The name of the module RMF is to invoke to format the header lines and data line(s) for the report.

RBLEN = length
The length of the relocate block generated by the data gatherer for each line in the report.

RPTTYP = {R|T}
The type of report. T indicates a table report; R indicates a row report.

TITLE = 'title'
An optional report title. The title specified appears in the menu frame for a display session. The title must be enclosed in single quotes. Use two quotes to represent any quote used in the title. The title can contain up to 50 printable characters. However, a maximum of 35 characters can be printed or displayed; therefore, a title longer than 35 characters will be truncated to fit into the menu frame.

DGTEXT = 'dgdefaults'
The default operands that are passed to the data-gathering routine for the report. This field is optional; it is used when the report requires operands. The text must be enclosed in single quotes, and the maximum length of the text is 32 characters. Any characters are valid between the quotes. Use two quotes to represent any quote used in the text. When more than 32 characters are specified, the text is truncated.

DRTEXT = 'drdefaults'
The default operands that are passed to the data-reporting routine for the report. This field is optional; it is used when the report requires operands. The text must be enclosed in single quotes. Use two single quotes to represent any quotes used in the text. Any characters are valid between the quotes. When more than 32 characters are specified, the text is truncated.

MAXRBS = nnn
The maximum number of relocate blocks. This number is equivalent to the maximum number of data lines in the report. The field is optional; when it is omitted, the default is 1 when RPTTYP = R is specified, indicating a row report. When RPTTYP = T is specified, indicating a table report, the field defaults to zero; however, enough storage is provided to allow a relocate block for each address space possible in the system. The maximum value possible for MAXRBS is 32,767.

FBLEN = len
The total length of all data control sections of the SMF record. The default value is 0.

Example

The following example shows how to use the PICTURE macro to add a menu item to ERBFMENU. The menu item for the report is ANL, the data gatherer is ANLDG, the PF key is 23, the data reporter is ANLRP, the length of the relocate block is 32, the length of all data control sections is 0, and the report is a table report. The title of the report is USER ANALYSIS, the default operands for the gatherer and the reporter are 1,1,1. The maximum number of relocate blocks is 128.

```
ANLPIC PICTURE ID=ANL,GATHER=ANLDG,PFK=23,REPORT=ANLRP,RBLEN=32,FBLEN=0,
               RPTTYP=T,TITLE='USER ANALYSIS',DGTEXT='1,1,1',
               DRTEXT='1,1,1'MAXRBS=128
```


TSO Terminal User Authorization

All the data collected and reported by RMF during a Monitor II TSO display session is obtained from commonly addressable storage that is not fetch protected. However, if your installation wants to limit the use of the command that starts an RMF Monitor II (RMFMON) session under TSO, one method available is to replace the RMF control section with your own module. For Monitor II you replace the control section ERBTSOCK. Your routine will then be invoked as part of the RMF response to the RMFMON command.

ERBTSOCK (Monitor II) runs in problem state with a key of 8. When this control section gets control, register 1 points to a two-word address list. The first address points to the seven-byte userid of the user who has issued the RMFMON command. The second word points to the PSCB. Figure 6-13 illustrates the input parameter structure.

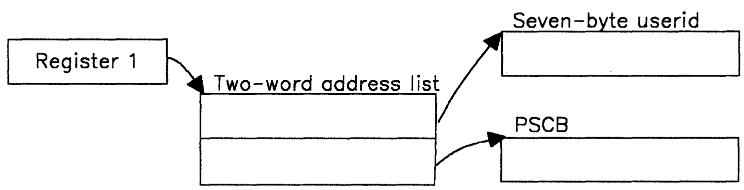

Figure 6-13. ERBTSOCK Input Parameter Structure

The module that you code to replace ERBTSOCK must be reenterable. It receives control by a BALR instruction and must save the registers when it receives control and restore the registers when it returns control. Register 13 contains the address of the register save area; register 14 contains the return address; and register 15 contains the entry address.

The processing your module performs depends on the method you choose to validate the user. Possible methods include issuing a RACHECK, prompting the user for a password, or checking the userid against a list of valid userids. Information on the TSO services available to perform these functions, such as TGET or TPUT, can be found in *TSO Guide to Writing a Terminal Monitor Program or a Command Processor*.

You can also use the PSCB bits defined for user use. This field (PSCBATR2 in the PSCB) comes from the UADS and can be updated by the USERDATA keyword of the ADD and CHANGE subcommands of the ACCOUNT command. See *System Programming Library: TSO* for further information.

When your routine has completed its processing, set a return code of 0 in register 15 to indicate to RMF that the user is authorized to issue RMFMON. Set a return code of 4 in register 15 to indicate to RMF that the user is not authorized to issue RMFMON. In response to this return code, RMF will display a message at the display station. No session will be started. After setting the appropriate return code, return control by branching on the contents of register 14.

For the Monitor II TSO session the user authorization exit routine (ERBTSOCK) is part of the RMF load module that contains the RMFMON command. This module resides in SYS1.LINKLIB as load module RMFMON; its entry point is ERBMFTSO. Before your authorization routine can execute, you must link edit it with RMFMON; the control statements required are:

```
(ERBTSOCK object deck)
INCLUDE ddname(RMFMON)
ENTRY ERBMFTSO
NAME RMFMON(R)
```


Chapter 7. SMF Records

Each SMF record contains information similar to the contents of the corresponding formatted report. For each system activity that you select, RMF collects data and formats an SMF record to hold the data it collects.

Some totals, averages, and percentages are not explicitly contained in the SMF records, but are calculated from the SMF data. For elaboration of particular fields, see the descriptions of the corresponding fields in the printed report descriptions in Chapters 3 and 4. Also, each SMF record produced by RMF is described in *System Management Facilities (SMF)*.

The record types and the corresponding RMF measurement activities are:

- Record Type 70 – CPU activity
- Record Type 71 – paging activity
- Record Type 72 – workload activity
- Record Type 73 – channel path activity
- Record Type 74 – device activity
- Record Type 75 – page/swap data set activity
- Record Type 76 – trace activity
- Record Type 77 – enqueue activity
- Record Type 78 – Monitor I extension
- Record Type 79 – Monitor II activity

Record type 78 and record type 79 have sub-types.

Record type 78 has the following sub-types:

- Sub-type 1 – I/O queuing activity for the 308x, 908x, and 4381 processors
- Sub-type 2 – Virtual storage activity
- Sub-type 3 – I/O queuing activity for the 3090 processor

Record type 79 has the following sub-types:

- Sub-type 1 – address space state data
- Sub-type 2 – address space resource data
- Sub-type 3 – real storage/processor/SRM
- Sub-type 4 – paging
- Sub-type 5 – address space SRM data
- Sub-type 6 – reserve data
- Sub-type 7 – enqueue contention data
- Sub-type 8 – transaction activity data
- Sub-type 9 – device activity
- Sub-type 10 – domain activity
- Sub-type 11 – paging activity
- Sub-type 12 – channel path activity
- Sub-type 13 – I/O queuing activity for the 308x, 908x, and 4381 processors
- Sub-type 14 – I/O queuing activity for the 3090 processor

SMF Record Format

Depending on the feedback options you select, RMF can write the SMF records to the SMF data set, use the data in the record to generate a printed report, or both. Regardless of the options you select, the format of the SMF record is the same.

Each SMF record that RMF generates consists of the following sections:

1. **SMF common header**, which identifies the record length, the record type, the time and date, the SMF system identifier, the subsystem identifier (always RMF), and the record subtype (if required). It also describes the other sections in the record. Each section is identified by its offset, the length of the section, and the number of such sections in the record. These offset/length/number triplet pointers define the structure of the rest of the record.
2. **RMF product section**, which includes information such as the RMF version number, the start time of the interval, the length of the interval, and the length of the sampling cycle. The RMF product section is the same in all records.
3. **Control section**, which contains general one-time data for RMF to use to produce any requested report. The contents of the section depend on the record type. Some records do not require a control section, while others require more than one.
4. **Data section**, which includes the specific data gathered during the interval. The format and the number of the data sections depend on the record type and the data collected. For example, there would be one data section for each device included in the type 74 record, I/O device activity.

With this format, the SMF records that RMF generates can change to incorporate any new or modified data without creating incompatibilities. The key factors in allowing for compatible change are the grouping of similar data in one section and the use of the offset/length/number triplet pointers to access the data stored in each section. Figure 7-1 shows the general format of the SMF records that RMF generates. The figure shows both the pointer structure and the storage layout for the sections.

Also, you can access fields in the SMF common header and the RMF product section by either a general name or a specific name. For example, you can access the interval start time in a type 70 record by either its general name (SMFIST) or its specific name (SMF70IST). Thus, code that processes all records can use the general name while code that processes only a specific record type can use the specific name.

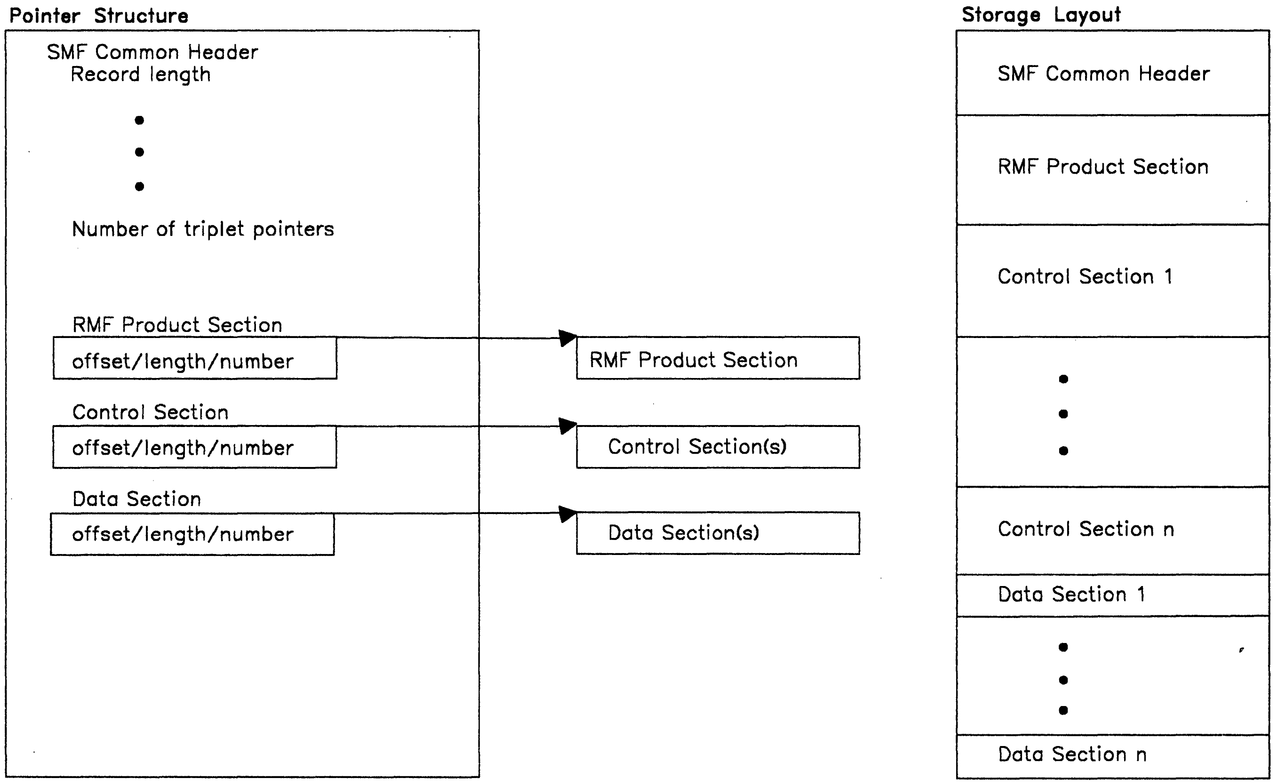

Figure 7-1. SMF Record Format

If your installation has existing data reduction programs that use SMF record input, check the SMF record formats carefully to determine what changes are required. Note that using the SMF record mapping macro instructions supplied by RMF is the most flexible way to access the contents of the SMF records your programs require. When you use the mapping macros, usually only a re-assembly of your program is required to incorporate changes to the record format.

The SMF record mapping macro instruction is ERBSMFR. Its format is:

```
ERBSMFR(nn1[,nn....1])
```

where nn identifies the type(s) of the SMF record(s) you want to map. Note that the parentheses are required only when two or more SMF record types are specified.

If you specify ERBSMF, the macro generates a mapping of the SMF common header and the RMF product section using only the general names.

The mapping macros reside in SYS1.RMFMAC01.

Because RMF can generate spanned SMF records – particularly when I/O device activity is measured – correct DCB parameters are important. Do not override the DCB parameters in the data set label by specifying DCB parameters on JCL statements. However, when using unlabeled tape the JCL describing an input SMF record data set should specify RECFM=VBS and a logical record length (LRECL) that is at least equal to the length of the longest record.

Archived Performance Data

Many installations find it useful to archive the performance data collected in the SMF records RMF produces and use this data, for example, to study trends or to evaluate the impact of a system change. Because of system changes and/or RMF changes, the archived data recorded by various versions or releases of RMF is not always the same.

The Version 3 post processor, however, can process records produced by some previous versions or releases of RMF. The post processor determines the RMF version number of each SMF record in the input stream. The RMF level appears in field name SMFxxMFV, where xx is the record number. If the field contains one of the following values, the post processor rejects the record:

- X'F0F1' for an SMF record produced by MF/1
- X'F0F2' for an SMF record produced by RMF Version 1
- X'F0F03' for an SMF record produced by RMF Version 2 Release 1 and RMF Version 2 Release 2
- X'F0F4' for an SMF record produced by RMF Version 2 Release 2 when MVS/System Extensions Release 1 is installed

If the field contains one of the following values, the post processor accepts the record and uses the data it contains to produce reports:

- X'F0F5' for an SMF record produced by RMF Version 2 Release 2 when MVS/System Extensions Release 2 is installed or for an SMF record produced by RMF Version 2 Release 3
- X'F0F6' for an SMF record produced by RMF Version 2 Release 3 when either MVS/System Product-JES2 Release 1 Enhancements (5740-XYS) or MVS/System Product-JES3 Release 1 Enhancements (5740-XYN) is installed.
- X'F0F7' for an SMF record produced by RMF Version 2 Release 4
- X'F0F8' for an SMF record produced by RMF Version 2 Release 4 Enhancements
- X'F3F1' for an SMF record produced by RMF Version 3 Release 1
- X'F3F2' for an SMF record produced by RMF Version 3 Release 2
- X'321F' for an SMF record produced by RMF Version 3 Release 2 Modification Level 1
- X'330F' for an SMF record produced by RMF Version 3 Release 3 Modification Level 0
- X'350F' for an SMF record produced by RMF Version 3 Release 5 Modification Level 0
- X'351F' for an SMF record produced by RMF Version 3 Release 5 Modification Level 1

When the version number indicates that the record was produced by an earlier version or release of RMF (SMFxxMFV contains (X'F0F5' through X'F0F8', X'F3F1', X'F3F2', or X'F321'), the post processor converts the record to the current RMF format. A converted record, however, is not exactly the same as a current record. The major differences are:

- Fields that contain data that only Version 3 collects contains blanks or zeroes in the converted record.
- Fields that contain data that Version 3 does not collect are omitted.
- The converted record contains a flag that indicates that it is a converted record, but RMF does preserve the original record version number.

Because of these differences, reports based on converted records will also differ from reports based on RMF Version 3 records. The major differences are:

- Fields that only Version 3 reports contain blanks or zeroes.
- Fields that Version 3 does not report are omitted.
- The report heading includes the original record version number and the current version number in the form PRT VERSION x.x.x TO y.y.y CONVERTED, where x.x.x is the version number at the time the data was gathered, and y.y.y is the currently installed RMF version at the time of the reporting.

When you compare data from a report based on current records with data from a report based on converted records, consider the differences between MVS/XA and MVS/370. For example, data values that might indicate constraints for an MVS/370 system might not signal problems for an MVS/XA system.

Also note that, particularly when the data pertains to I/O operations, some fields in converted records contain data that has no Version 3 equivalent but that is similar to data that Version 3 reports. In these cases, the post processor converts the data but reports it under the Version 3 heading. The reports that contain such converted data are channel path activity (Monitor I and Monitor II), I/O device activity (Monitor I and Monitor II), and paging activity (Monitor I).

Channel Path Activity: The channel number in the Version 3 report contains 2 digits. When the report comes from converted Version 2 data, the first digit is the CSID (channel set identifier) or CPUID (processor identifier). The second digit is the channel number.

I/O Device Activity: The following list identifies the Version 3 heading for each field that can contain converted data and the meaning of the converted data when the report comes from Version 2 records:

Version 3 Heading	Version 2 Contents
DEV NUM	Device address.
LCU	Logical channel number. Data is summarized by logical channel.
DEVICE ACTIVITY RATE	Start I/O rate.
AVE RESP TIME AVE IOSQ TIME	Total device active time is part of the calculation for these two fields. It is calculated from the Version 2 data in the same way as total service time was for the Version 2 report. That is: $\frac{\text{\#SAMPLES DEV BUSY}}{\text{\#SAMPLES}} \times \text{INTERVAL LENGTH}$
% DEV UTIL	Data reported for Version 2 under % DEV BUSY. Device busy represents the time during the interval when the device could not be used to service another request from this system. (Version 3 device utilization represents the time when the device could not be used to service a request from another system.)

Paging Activity: In the frame count section, some Version 2 records contain data on pageable frames and some contain data on total frames. (The type of data depends on whether the report was generated before or after your installation applied the fix for APAR OZ40661.) When the record contains pageable frame counts, the post processor combines the data, where appropriate, with fixed frame counts so that the converted records contain total frame counts. It is not appropriate to combine maximum or minimum frame counts because the result of the calculation would be invalid. It is appropriate, however, to compute average values by combining pageable and fixed frame counts. Therefore, a converted records reports average LPA, CSA, and private total frame counts but reports zeroes in the minimum and maximum fields for LPA, CSA, and private total frame counts.

In a report based on converted data, the count of pages fixed below the 16 megabyte line is set to the total count of fixed pages. (In MVS/370, all fixed pages resided below the 16 megabyte line.)

Printing SMF Records

You might occasionally find it necessary to print the SMF records RMF produces. Printed records are useful, for example, when designing and implementing a user-written record processing program or when diagnosing problems with RMF reports. A sample of the JCL needed to print RMF records follows. The first step (SELECT) limits the amount of output to the record types or time frames that you need. If you want to print the entire data set, use only the second step (PRINT). These JCL statements and SMF dump parameters select and print SMF record types 70 through 79 that were written from 10:00 AM until noon on April 26, 1983.

```
//SELECT     EXEC  PGM=IFASMFDP
//SYSPRINT   DD    SYSOUT=A
//IN         DD    DSN=data set containing SMF records
//OUT        DD    DSN=&&RMFREC,DISP=(NEW,PASS),UNIT=SYSDA
//SYSIN      DD    *
  INDD(IN,OPTIONS(DUMP))
  OUTDD(OUT,TYPE(70:79))
  START(1000)
  END(1200)
  DATE(83116,83116)
/*
//PRINT      EXEC  PGM=IDCAMS
//SYSPRINT   DD    SYSOUT=A
//RMFREC     DD    DSN=&&RMFREC,DISP=(OLD,PASS)
//SYSIN      DD    *
  PRINT   INFILE(RMFREC)
/*
```

You can also select and print just one SMF record type by specifying

```
OUTDD(OUT,TYPE(74))
```

on the SYSIN DD statement in the SELECT step.

Note: If you choose to print the entire data set and execute only the second step, specify

```
DSN=data set containing records
```

on the RMFREC DD statement.

MVS/Extended Architecture System Programming Library: System Management Facilities (SMF) contains more information on IFASMFDP dump program. *MVS/XA Access Methods Services Reference for the Integrated Catalog Facility, GC26-4019* and *MVS/XA Access Methods Services Reference for VSAM Catalogs, GC26-4075* contains more information about IDCAMS.

Because you do not specify the format on the PRINT statement, the format defaults to DUMP. The records are printed in a dump format. Figure 7-2 is an example of the SMF record dump format. The offsets are in the left column, and the right side of the dump contains a printable section to help find the fields of interest. Note that the PRINT utility does not include the record length and segment descriptor fields in its output. As a result, a field shown at offset 4 in an SMF record in *MVS/XA SPL:SMF* appears at offset 0 in the formatted dump. You must adjust subsequent offsets accordingly to refer back and forth from the formatted dump to the printed SMF records in the *MVS/XA SPL:SMF* manual.


```
IDCAMS    SYSTEM SERVICES                             TIME: 09:53:41          04/20/87        PAGE 43

LISTING OF DATA SET _SYS83110.T095324.RA000.D09BAB1A.BARB

RECORD SEQUENCE NUMBER - 1
000000  02020036 54790083 110FE2D7 F2F1                                        *.........SP21

RECORD SEQUENCE NUMBER - 1
000000 8646002B F2720083 110FE2D7 F2F1D9D4   C6400001 00040000 0000003C 00280001  *....2.....SP21RMF .............*
000020 00000064 00030001 00000067 00180002   00000097 02080001 F3F2D9D4 C6404040  *.......................32RMF   *
000040 40400070 131F0083 110F5829 510F0000   00000D85 00000000 F0F3F840 0001000F  *.......................038 ....*
000060 30810300 00B211B  919A0000 00010002   00980000 08BE6B00 002B2A00 000B42DF  *...............................*
000080 53120000 02010022 00009800 009D1A00   00000000 00000000 00106900 00000000  * ..............................*
0000A0 000B0000 00019800 00007000 00000000   00000000 00000000 00000000 00000000  *...............................*
0000C0 00000000 00000000 00000000 00000000   00000000 00000000 00000000 09001200  *...............................*
0000E0 00005300 00000000 00000000 00000000   00000000 00000000 00000000 00756000  *...............................*


RECORD SEQUENCE NUMBER - 2
000000 8646002B F2720083 110FE2D7 F2F1D9D4   C6400001 00040000 0000003C 00280001  *....2.....SP21RMF .............*
000020 00000064 00030001 00000067 00180002   00000097 02080001 F3F2D9D4 C6404040  *.......................32RMF   *
000040 40400070 131F0083 110F5829 510F0000   00000D85 00000000 F0F3F840 0001000F  *.......................038 ....*
000060 30810300 00B211B  919A0000 00010002   00980000 08BE6B00 002B2A00 000B42DF  *...............................*
000080 53120000 02010022 00009800 009D1A00   00000000 00000000 00106900 0)000000  * ..............................*

RECORD SEQUENCE NUMBER - 3
000000 8646002B F2720083 110FE2D7 F2F1D9D4   C6400001 00040000 0000003C 00280001  *....2.....SP21RMF .............*
000020 00000064 00030001 00000067 00180002   00000097 02080001 F3F2D9D4 C6404040  *.......................32RMF   *
000040 40400070 131F0083 110F5829 510F0000   00000D85 00000000 F0F3F840 0001000F  *.......................038 ....*
000060 30810300 00B211B  919A0000 00010002   00980000 08BE6B00 002B2A00 000B42DF  *...............................*
000080 53120000 02010022 00009800 009D1A00   00000000 00000000 00106900 0)000000  * ..............................*
```

Figure 7-2. Dump Format of SMF Record

Chapter 8. Monitor I Fixed Storage Requirements

Fixed Storage Estimates (by options) for Monitor I

The following is a table you can use to determine the fixed storage requirements for the Monitor I session, according to the options you specify. This table has been filled in for reference use. A blank table follows this one for your own use.

Note: Common storage area (CSA) estimates for Monitor I fixed storage are not included.

Variables		Formulas/Estimates		
If channel (CHAN) activity measurement is selected				
A1	– total SQA requirement	140	140	140
A2	– total temporarily fixed address space for CHAN	A2 = total	1,024	1,024
If processor (CPU) activity measurement is selected				
B1	– total SQA requirement		400	400
B2	– PLPA (fixed)		2,064	2,064
B3	– basic requirement		3,388	
B4	– number of processors (assume 1)	16*B4	16	
B5	– total temporarily fixed address space for CPU	B5 = total	3,404	3,404
If device (DEVICE) activity measurement is selected				
C1	– basic SQA requirement	150	150	
C2	– number of devices sysgened in device classes being measured (assume 275)	68*C2	18,700	
C3	– total SQA for DEVICE	C3 = total	18,850	18,850
C4	– PLPA code (fixed)		744	744
C5	– address space code - basic requirement		6064	
C6	– number of device classes being measured (assume 2)	10*C6	20	
C7	– number of devices sysgened in device classes being measured (assume 275)	72*C7	19,800	
C8	– total temporarily fixed address space code for DEVICE	= total	25,884	25,884
If paging (PAGING) activity measurement is selected				
D1	– total SQA required	486	486	486
D2	– PLPA (fixed)	1,376	1,376	1,376
D3	– total temporarily fixed address space code requirement	4.8	4,880	4,880
SUBTOTAL 1				59,252

Variables		Formulas/Estimates		
If Page/Swap data set (PAGESP) activity measurement is selected				
E1	– basic SQA requirement	24	24	
E2	– maximum number of page data sets which could be varied online for this IPL (up to 64) – assume 5	52*E2	260	
E3	– maximum number of swap data sets which could be varied online for this IPL (up to 25) – assume 2	52*E3	104	
E4	– total SQA for PAGESP	E4 = total	388	388
E5	– PLPA code (fixed)	1K	1,120	1,120
E6	– address space code basic requirement	3.2K	3,191	
E7	– maximum number of page and swap data sets that could be varied online for this IPL (up to 89) – assume 7	164*E7	1,148	
E8	– total temporarily fixed address space code requirement	E8 = total	4,339	4,339
If ASM/RSM/SRM trace (TRACE) activity measurement is selected				
F1	– basic SQA requirement per name	100	100	
F2	– set size per name (depends on options chosen and the name's field size):	2 to 40	20	

Option Requested	Length of Sampled Data ≤2 bytes	>2 bytes
Minimum	2	4
Maximum	2	4
End	2	4
Average	6	8
Standard Deviation	14	20
All	20	32

F3	– samples per set – default is 60 (can be zapped in ERBMFTTB)	60	60	
F4	– number of samples (interval time ÷ cycle time) default interval = 15 minutes default cycle = 250 milliseconds } 3600 samples	6 to 72000	3,600	
F5	– total SQA for a name $\left(\frac{F4}{F3} * F2\right) + F1$ F5 must be calculated for each name.	$\left(\frac{F4}{F3} * F2\right) + F1$	1,300	
F6	– total SQA for all names traced = summation of space for each name	Σ F5	6,500	6,500
F7	– PLPA code (fixed)	1808	1,808	1,808
F8	– total temporarily fixed address space for TRACE data	2652 + F6	9,152	9,152
SUBTOTAL 2				23,307

Notes:

1. Requesting one of the 5 Domain Names will expand the number of names traced.
2. Discretion must be used in selecting the list of names to trace and the options. If the entire set of names and options is requested with the shortest cycle and the longest interval, an excessive amount of storage is required. You should make a rough estimate of your requirements, how much storage you will need, and what effect this will have on the system performance. By not requesting standard deviation, you will reduce storage requirements significantly (See the table with F2.)

Variables	Formulas/Estimates		
If workload (WKLD) activity measurement is selected			
G1 – basic SQA requirement	64	64	
G2 – number of performance numbers in Installation Performance Specification (IPS) (maximum = 255)	4*G2	40	
G3 – number of performance group periods in IPS (maximum = 2040)	44*G3	880	
G4 – total SQA for WKLD	G4 = total	984	984
G5 – total temporarily fixed address space requirements	2.9K	2,908	
G6 – highest performance group number in IPS (maximum = 255)	4*G6	24	
G7 – number of performance group periods in IPS	44* G7	880	
G8 – total temporarily fixed address space code for WKLD	G8 = total	3,812	3,812
If enqueue (ENQ) activity measurement is selected			
H1 – basic SQA requirement for ENQ on a specified major name	.4K	408	
H2 – basic SQA requirement for ENQ on a specified major name and minor name	2K	2,048	
H3 – basic SQA requirement for ENQ on all serially-reusable resources	4K	4,096	
H4 – total SQA requirement (assume H2 value)	H4 = total	2,048	2,048
H5 – PLPA (fixed)	4.6K	4,656	4,656
H6 – total temporarily fixed address space for ENQ	5.7K	5,734	5,734
If I/O Queuing (IOQ) activity measurement is selected on a 308x, 908x, or 4381 processor			
I1 – basic SQA requirement	192	192	
I2 – number of LCUs (assume 50)	144*I2	7,200	
I3 – total SQA for IOQ	I3 = total	7,392	7,392
I4 – PLPA (fixed)		900	900
I5 – total temporarily fixed address space code for IOQ	4K	4,096	4,096
If I/O Queuing (IOQ) activity measurement is selected on a 3090 processor			
I1 – basic SQA requirement	324	324	
I2 – number of LCUs (assume 50)	230*I2	11,400	
I3 – total SQA for IOQ	I3 = total	11,724	11,724
I4 – PLPA (fixed)		900	900
I5 – total temporarily fixed address space code for IOQ	4K	4,096	4,096
If Virtual Storage (VSTOR) activity measurement is selected			
J1 – basic SQA requirement	11.8K	11,796	
J2 – number of monitored jobs	44+5524*J2	5,568	
J3 – total SQA for VSTOR	J3 = total	17,364	17,364
J4 – PLPA (fixed)	4.9K	4,908	4,908
J5 – total temporarily fixed address space code for VSTOR	4620+3872*J2	8,492	8,492
SUBTOTAL 3 (with I/O Queuing on a 308x, 908x, or 4381 processor)			60,386
SUBTOTAL 3 (with I/O Queuing on a 3090 processor)			64,718
SUBTOTAL 2			23,307
SUBTOTAL 1			59,252
TOTAL FIXED RMF CODE (with I/O Queuing on a 308x, 908x, or 4381 processor)			142,945
TOTAL FIXED RMF CODE (with I/O Queuing on a 3090 processor)			147,227

Monitor I Fixed Storage Estimates (by storage area)

Storage Area	Formula	Estimate
SQA for RMF		
Basic requirement	8,332	8,332
CPU (if specified)	B2	400
CHAN (if specified)	A1	140
DEVICE (if specified)	C3	18,850
PAGING (if specified)	D1	486
WKLD (if specified)	G4	984
PAGESP (if specified)	E4	388
TRACE (if specified)	F6	6,500
ENQ (if specified)	H4	2,048
IOQ (if specified for a 308x, 908x, or 4381 processor)	I3	7,392
IOQ (if specified for a 3090 processor)	I3	11,724
VSTOR (if specified)	J3	17,364
	total[1]	62,884
	total[2]	67,216
Fixed PLPA for RMF		
Basic requirement		1,368
CPU (if specified)	B2	2,064
CHAN (if specified)		0
DEVICE (if specified)	C4	744
PAGING (if specified)	D2	1,376
WKLD (if specified)		0
PAGESP (if specified)	E5	1,120
TRACE (if specified)	F7	1,808
ENQ (if specified)	H5	4,656
IOQ (if specified for 308x, 908x, or 4381	I4	900
IOQ (if specified for a 3090 processor)	I4	900
VSTOR (if specified)	J4	4,908
	total	19,844
Temporarily Fixed Private Area Storage for RMF		
Basic requirement	.7K	718
CPU (if specified)	B5	3,404
CHAN (if specified)	A2	1,024
DEVICE (if specified)	C8	25,884
WKLD (if specified)	G9	3,812
PAGING (if specified)	D3	4,880
PAGESP (if specified)	E8	4,339
TRACE (if specified)	F8	9,152
ENQ (if specified)	5.7K	5,734
IOQ (if specified)	I5	4,096
VSTOR (if specified)	J5	8,492
	total	69,523

[1] for a 308x, 908x, or 4381 processor

[2] for a 3090 processor

This form is provided for your use

Variables	Formulas/Estimates		
If channel (CHAN) activity measurement is selected			
A1 – total SQA requirement	140		
A2 – total temporarily fixed address space for CHAN	A2 = total		
If processor (CPU) activity measurement is selected			
B1 – total SQA requirement			
B2 – PLPA (fixed)			
B3 – basic requirement			
B4 – number of processors (assume 1)	16*B4		
B5 – total temporarily fixed address space for CPU	B5 = total		
If device (DEVICE) activity measurement is selected			
C1 – basic SQA requirement	150		
C2 – number of devices sysgened in device classes being measured (assume 275)	58*C2		
C3 – total SQA for DEVICE	C3 = total		
C4 – PLPA code (fixed)			
C5 – address space code - basic requirement			
C6 – number of device classes being measured (assume 2)	10*C6		
C7 – number of devices sysgened in device classes being measured (assume 275)	72*C7		
C8 – total temporarily fixed address space code for DEVICE	= total		
If paging (PAGING) activity measurement is selected			
D1 – total SQA required	486		
D2 – PLPA (fixed)	1,376		
D3 – total temporarily fixed address space code requirement	4.8		
SUBTOTAL 1			

Variables	Formulas/Estimates		
If Page/Swap data set (PAGESP) activity measurement is selected			
E1 – basic SQA requirement	24		
E2 – maximum number of page data sets which could be varied online for this IPL (up to 64) – assume 5	52*E2		
E3 – maximum number of swap data sets which could be varied online for this IPL (up to 25) – assume 2	52*E3		
E4 – total SQA for PAGESP	E4 = total		
E5 – PLPA code (fixed)	1K		
E6 – address space code basic requirement	3.2K		
E7 – maximum number of page and swap data sets that could be varied online for this IPL (up to 89) – assume 7	164*E7		
E8 – total temporarily fixed address space code requirement	E8 = total		
If ASM/RSM/SRM trace (TRACE) activity measurement is selected			
F1 – basic SQA requirement per name	100		
F2 – set size per name (depends on options chosen and the name's field size):	2 to 40		
F3 – samples per set – default is 60 (can be zapped in ERBMFTTB)	60		
F4 – number of samples (interval time ÷ cycle time) default interval = 15 minutes, default cycle = 250 milliseconds } 3600 samples	6 to 72000		
F5 – total SQA for a name $\left(\frac{F4}{F3} * F2\right) + F1$ F5 must be calculated for each name.	$\left(\frac{F4}{F3} * F2\right) + F1$		
F6 – total SQA for all names traced = summation of space for each name	Σ F5		
F7 – PLPA code (fixed)	1808		
F8 – total temporarily fixed address space for TRACE data	2652 + F6		
SUBTOTAL 2			

Table with F2:

Option Requested	Length of Sampled Data ≤2 bytes	>2 bytes
Minimum	2	4
Maximum	2	4
End	2	4
Average	6	8
Standard Deviation	14	20
All	20	32

Notes:

1. Requesting one of the 5 Domain Names will expand the number of names traced.
2. Discretion must be used in selecting the list of names to trace and the options. If the entire set of names and options is requested with the shortest cycle and the longest interval, an excessive amount of storage is required. You should make a rough estimate of your requirements, how much storage you will need, and what effect this will have on the system performance. By not requesting standard deviation, you will reduce storage requirements significantly (See the table with F2.)

Variables		Formulas/Estimates		
If workload (WKLD) activity measurement is selected				
G1	– basic SQA requirement	64		
G2	– number of performance numbers in Installation Performance Specification (IPS) (maximum = 255)	4*G2		
G3	– number of performance group periods in IPS (maximum = 2040)	44*G3		
G4	– total SQA for WKLD	G4 = total		
G5	– total temporarily fixed address space requirements	2.9K		
G6	– highest performance group number in IPS (maximum = 255)	4*G6		
G7	– number of performance group periods in IPS	44* G7		
G8	– total temporarily fixed address space code for WKLD	G8 = total		
If enqueue (ENQ) activity measurement is selected				
H1	– basic SQA requirement for ENQ on a specified major name	.4K		
H2	– basic SQA requirement for ENQ on a specified major name and minor name	2K		
H3	– basic SQA requirement for ENQ on all serially reusable resources	4K		
H4	– total SQA requirement (assume H2 value)	H4 = total		
H5	– PLPA (fixed)	4.6K		
H6	– total temporarily fixed address space for ENQ	5.7K		
If I/O Queuing (IOQ) activity measurement is selected on a 308x, 908x, or 4381 processor				
I1	– basic SQA requirement	192		
I2	– number of LCUs (assume 50)	144*I2		
I3	– total SQA for IOQ	I3 = total		
I4	– PLPA (fixed)			
I5	– total temporarily fixed address space code for IOQ	4K		
If I/O Queuing (IOQ) activity measurement is selected on a 3090 processor				
I1	– basic SQA requirement	324		
I2	– number of LCUs (assume 50)	230*I2		
I3	– total SQA for IOQ	I3 = total		
I4	– PLPA (fixed)			
I5	– total temporarily fixed address space code for IOQ	4K		
If Virtual Storage (VSTOR) activity measurement is selected				
J1	– basic SQA requirement	11.8K		
J2	– number of monitored jobs	44+5524*J2		
J3	– total SQA for VSTOR	J3 = total		
J4	– PLPA (fixed)	4.9K		
J5	– total temporarily fixed address space code for VSTOR	4620+3872*J2		
SUBTOTAL 3 (with I/O Queuing on a 308x, 908x, or 4381 processor)				
SUBTOTAL 3 (with I/O Queuing on a 3090 processor)				
SUBTOTAL 2				
SUBTOTAL 1				
TOTAL FIXED RMF CODE (with I/O Queuing on a 308x, 908x, or 4381 processor)				
TOTAL FIXED RMF CODE (with I/O Queuing on a 3090 processor)				

Monitor I Fixed Storage Estimates (by storage area)

Storage Area	Formula	Estimate
SQA for RMF		
Basic requirement	8,332	
CPU (if specified)	B2	
CHAN (if specified)	A1	
DEVICE (if specified)	C3	
PAGING (if specified)	D1	
WKLD (if specified)	G4	
PAGESP (if specified)	E4	
TRACE (if specified)	F6	
ENQ (if specified)	H4	
IOQ (if specified for a 308x, 908x, or 4381 processor)	I3	
IOQ (if specified for a 3090 processor)	I3	
VSTOR (if specified)	J3	
	total[1]	
	total[2]	
Fixed PLPA for RMF		
Basic requirement		
CPU (if specified)	B2	
CHAN (if specified)		
DEVICE (if specified)	C4	
PAGING (if specified)	D2	
WKLD (if specified)		
PAGESP (if specified)	E5	
TRACE (if specified)	F7	
ENQ (if specified)	H5	
IOQ (if specified for 308x, 908x, or 4381	I4	
IOQ (if specified for a 3090 processor)	I4	
VSTOR (if specified)	J4	
	total	
Temporarily Fixed Private Area Storage for RMF		
Basic requirement	.7K	
CPU (if specified)	B5	
CHAN (if specified)	A2	
DEVICE (if specified)	C8	
WKLD (if specified)	G9	
PAGING (if specified)	D3	
PAGESP (if specified)	E8	
TRACE (if specified)	F8	
ENQ (if specified)	5.7K	
IOQ (if specified)	I5	
VSTOR (if specified)	J5	
	total	

[1] for a 308x, 908x, or 4381 processor

[2] for a 3090 processor

Chapter 9. User Completion Codes for Monitor I and II

For an explanation of the problem determination items, see *Message Library: System Codes.*

0301

Explanation: A return code other than 0 or 16 was returned by the TPUT SVC.

System Action: RMF either terminates the session or issues message ERB235A to allow the user to decide whether or not to continue the session.

Programmer Response: None.

Problem Determination: Items 4, 5b, 16, and 29.

1200

Explanation: An unexpected return code was received from the Monitor I session initialization SVC because the SVC tried to issue an ENQ macro instruction but the name (SYSZRBMF,ACTIVE) was already in use.

System Action: Monitor I session processing terminates.

Programmer Response: None.

Problem Determination: Items 2, 4, 5a, 16, and 29.

1201

Explanation: An unexpected return code was received from the Monitor I session data collection SVC.

System Action: Monitor I session processing terminates.

Programmer Response: None.

Problem Determination: Items 2, 4, 5a, 16, and 29.

1202

Explanation: An OPEN error occurred while opening the RMF programmer message data set that was to contain abnormal messages issued as a result of errors detected during RMF processing.

System Action: RMF processing terminates.

Programmer Response: None.

Problem Determination: Items 2, 4, 5b, 16, and 29.

1203

Explanation: An error occurred while attempting to dynamically allocate a SYSOUT data set.

System Action: If the error is associated with a report subtask, RMF measurement collection continues.

If the error is associated with the main RMF task, processing terminates.

Problem Determination: Items 4, 5a, 16, and 29.

1204

Explanation: A return code other than 0 was returned by the I/O configuration retrieve module ERBCNFGR.

System Action: RMF processing terminates.

Programmer Response: None.

Problem Determination: Items 2, 4, 5a, 16, 29.

1205

Explanation: During initialization of a Monitor I or Monitor II background session, a return code was returned by the system service routine IEEMB787 when accessing the data set containing the session options. The session options are located in the data set named on the IEFPARM DD statement in the RMF cataloged procedure; the data set is normally SYS1.PARMLIB.

System Action: The session is not initialized.

Programmer Response: None.

Problem Determination: Items 4, 5b, 16, 29.

1206

Explanation: While gathering data for the Monitor I Virtual Storage Report, the sampler received a return code of 8 or greater from the VSMLIST service. The meaning of these return codes is described in *SPL: System Macros and Facilities, Volume 2, GC28-1151.*

System Action: An SDUMP is produced.

1. If the failure occurred while Monitor I gathered data for the common area (indicated by ISSUER=ERBMFEVT in the SDUMP title), Monitor I terminates measurement collection. RMF attempts to reinstate Monitor I and, if successful, Monitor I continues to measure data under the control of the original options. If reinstatement fails, RMF terminates Monitor I.
2. If the failure occurred while Monitor I gathered data for a private area (indicated by ISSUER=ERBMFPVS in the SDUMP title), RMF attempts one retry. If the retry fails for the same private area, RMF no longer collects data for that private area. If there are two failures for two different private areas, RMF terminates Monitor I.

Programmer Response: If Monitor I terminates, restart it without the virtual storage report by specifying the NOVSTOR option.

Problem Determination: Items 2, 4, 18, 29, 33.

1207

Explanation: An RMF lock has not been obtained within the required number of attempts.

System Action: An SDUMP is produced.

1. If the failure occurred in ERBMFPVS, the lock used to serialize with ERBMFDVP remained unavailable for more than 10 seconds.
2. If the failure occurred in ERBMFDVP, the lock used to serialize with ERBMFPVS remained unavailable during 20,000 attempts by ERBMFDVP to obtain it.
3. If the failure occurred in IGX00022, the lock STGSLOCK used to serialize with ERBMFEVT remained unavailable during 20,000 attempts by IGX00022 to obtain it.
4. If the failure occurred in ERBMFEVT, the lock used to serialize with IGX00022 remained unavailable for more than 10 seconds.

Programmer Response: None.

Problem Determination: Items 4, 5b 16, 18, and 29.

1300

Explanation: An OPEN error was detected while opening a report SYSOUT data set.

System Action: No reports are formatted for this measurement interval. RMF measurement collection continues.

Programmer Response: None.

Problem Determination: Items 4, 5b, 16, 18, and 29.

1301

Explanation: The report subtask experienced an unexpected return code from one of the system services while formatting and printing reports.

System Action: No reports are formatted for this measurement interval. RMF measurement collection continues.

Programmer Response: None.

Problem Determination: Items 4, 5b, 16, 18, and 29.

1310

Explanation: An I/O error occurred while writing a report to the report SYSOUT data set.

System Action: The report subtask might be reinstated. If the reinstatement is successful, a complete set of reports is reformatted and printed for this interval. If the reinstatement is unsuccessful, the reports for the RMF interval are not completely formatted and all RMF measurement collection continues.

Programmer Response: None.

Problem Determination: Items 4, 5b, 16, 18, and 29.

1401

Explanation: An error code of 4, 16, 20, or more than 20 was returned by the TGET SVC.

System Action: RMF either terminates the session or issues message ERB235A to allow the user to decide whether or not to continue the session.

Programmer Response: None.

Problem Determination: Items 4, 5b, 16, and 29.

1402

Explanation: A return code of 8, indicating an invalid entry code parameter, was returned by a data gatherer or data reporter during a Monitor II session.

System Action: RMF either terminates the session or issues message ERB235A to allow the user to decide whether or not to continue the session.

Programmer Response: None, if the module causing the error is supplied by IBM. If the module causing the error is an installation-supplied routine, locate the source of the error and test the module again.

Problem Determination: Items 4, 5b, 16, and 29.

1403

Explanation: During a Monitor II session, RMF module ERBRMFPL -- the putline routine -- was called with an invalid input parameter. Either the text length was 0 or greater then 79, or the data type contained a value other than HD or DT.

System Action: RMF either terminates the session or issues message ERB235A to allow the user to decide whether or not to continue the session.

Programmer Response: None, if the module causing the error is supplied by IBM. If the module causing the error is an installation-supplied routine, locate the source of the invalid parameter to the putline routine and test the module again.

Problem Determination: Items 4, 5b, 16, and 29.

1404

Explanation: During a display session, RMF module ERBPUTSM was called to add a header or data line to the logical terminal buffer, but the buffer already contained the maximum number of header or data lines.

System Action: RMF either terminates the session or issues message ERB235A to allow the user to decide whether or not to continue the session.

Programmer Response: None, if the module causing the error is supplied by IBM. If the module causing the error is an installation-supplied routine, locate the source of the error and test the module again.

Problem Determination: Items 4, 5b, 16, and 29.

1405

Explanation: During a display session, RMF module ERBPUTSM was called to add a header or data line to the logical terminal buffer, but the length specified for the text was zero.

System Action: RMF either terminates the session or issues message ERB235A to allow the user to decide whether or not to continue the session.

Programmer Response: None, if the module causing the error is supplied by IBM. If the module causing the error is an installation-supplied routine, locate the source of the error and test the module again.

Problem Determination: Items 4, 5b, 16, and 29.

1410

Explanation: An error code of 4, 16, 20, or more than 20 was returned by the TGET SVC.

System Action: RMF either terminates the session or issues message ERB276I to allow the user to decide whether or not to continue the session.

Programmer Response: None.

Problem Determination: Items 4, 5b, 16, and 29.

1604

Explanation: During a Monitor II data gatherer session, RMF module ERB3GJS2 found a broken SJB queue.

System Action: RMF issues an SDUMP with user completion code 1604, issues a message ERB269I, indicating a permanent error condition, and terminates the JES2 data gatherer session.

Programmer Response: Make sure that the offset table ERB3GJ2E is valid for your JES2 release. Contact your systems programmer.

Problem Determination: Items 4, 5b, 16, and 29.

Index

A

D

E

F

G

H

I

J

K

L

M

N

Q

R

T

U

V

W

X

Z

Special Characters

MVS/XA
Resource Measurement
Facility (RMF)
Monitor I and II
Reference and User's
Guide

READER'S COMMENT FORM

LC28-1556-2

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Note: ***Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.***

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

____

____

____

____

What is your occupation? ______________________________

How do you use this publication? ______________________________

Number of latest Newsletter associated with this publication: ____________________

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

MVS/XA Resource Measurement Facility (RMF) Monitor I and II Reference and User's Guide

LC28-1556-2 S370-34

c u t a l o n g t h i s l i n e

Reader's Comment Form

Fold and Tape Please Do Not Staple Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department D58, Building 921-2
PO Box 950
Poughkeepsie, New York 12602-9935

Fold and Tape Please Do Not Staple Fold and Tape

Printed in U.S.A.

IBM®

MVS/XA
Resource Measurement
Facility (RMF)
Monitor I and II
Reference and User's
Guide

READER'S
COMMENT
FORM

LC28-1556-2

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Note: ***Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.***

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

What is your occupation? ______

How do you use this publication? ______

Number of latest Newsletter associated with this publication: ______

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

MVS/XA Resource Measurement Facility (RMF) Monitor I and II Reference and User's Guide

LC28-1556-2 S370-34

cut along this line

Reader's Comment Form

Fold and Tape — Please Do Not Staple — Fold and Tape

NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department D58, Building 921-2
PO Box 950
Poughkeepsie, New York 12602-9935

Fold and Tape — Please Do Not Staple — Fold and Tape

Printed in U.S.A.

IBM ®

MVS/XA
Resource Measurement
Facility (RMF)
Monitor I and II
Reference and User's
Guide

READER'S COMMENT FORM

LC28-1556-2

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Note: ***Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.***

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

____

____

____

____

What is your occupation? ______________________________

How do you use this publication? ______________________________

Number of latest Newsletter associated with this publication: ____________________

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

MVS/XA Resource Measurement Facility (RMF) Monitor I and II Reference and User's Guide

LC28-1556-2 S370-34

c u t a l o n g t h i s l i n e

Reader's Comment Form

Fold and Tape Please Do Not Staple Fold and Tape

NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department D58, Building 921-2
PO Box 950
Poughkeepsie, New York 12602-9935

Fold and Tape Please Do Not Staple Fold and Tape

Printed in U.S.A.

MVS/XA
Resource Measurement
Facility (RMF)
Monitor I and II
Reference and User's
Guide

READER'S
COMMENT
FORM

LC28-1556-2

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Note: ***Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.***

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

If you wish a reply, give your name, company, mailing address, and date:

What is your occupation? ____________________

How do you use this publication? ____________________

Number of latest Newsletter associated with this publication: ____________________

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

MVS/XA Resource Measurement Facility (RMF) Monitor I and II Reference and User's Guide

c u t a l o n g t h i s l i n e

LC28-1556-2 S370-34

Reader's Comment Form

Fold and Tape Please Do Not Staple Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department D58, Building 921-2
PO Box 950
Poughkeepsie, New York 12602-9935

Fold and Tape Please Do Not Staple Fold and Tape

Printed in U.S.A.

IBM ®

MVS/Extended Architecture Resource Measurement Facility (RMF) Monitor I and II Reference and User's Guide

LC28-1556-2

S370-34

Printed in U.S.A.

LC28-1556-02