z/OS
Version 2 Release 4

Resource Measurement Facility Data Gatherer User's Guide

SC27-4934-40

Note

Before using this information and the product it supports, read the information in “Notices” on page 93.

This edition applies to Version 2 Release 4 of z/OS (5650-ZOS) and to all subsequent releases and modifications until otherwise indicated in new editions.

Last updated: 2019-11-20

Contents

Figures

Tables

About this document

The Resource Measurement Facility (RMF) is the strategic IBM® product for performance management in a z/OS host environment.

This document describes RMF, what it can do, and how to use RMF sessions. For information about analyzing the various reports that RMF produces, see *z/OS RMF Report Analysis*.

Who should use this document

This document is intended for use by:

- System administrators and programmers responsible for installing RMF and modifying its functions
- Performance analysts responsible for measuring and improving system performance,
- System operators

Because RMF is a product for measuring system performance of a z/OS system, this document assumes that the reader has extensive knowledge of the z/OS system.

z/OS information

This information explains how z/OS references information in other documents and on the web.

When possible, this information uses cross document links that go directly to the topic in reference using shortened versions of the document title. For complete titles and order numbers of the documents for all products that are part of z/OS, see *z/OS Information Roadmap*.

To find the complete z/OS® library, go to IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSLTBW/welcome).

How to read syntax diagrams

This section describes how to read syntax diagrams. It defines syntax diagram symbols, items that may be contained within the diagrams (keywords, variables, delimiters, operators, fragment references, operands) and provides syntax examples that contain these items.

Syntax diagrams pictorially display the order and parts (options and arguments) that comprise a command statement. They are read from left to right and from top to bottom, following the main path of the horizontal line.

For users accessing the Information Center using a screen reader, syntax diagrams are provided in dotted decimal format.

Symbols

The following symbols may be displayed in syntax diagrams:

Symbol
Definition

►►—
Indicates the beginning of the syntax diagram.

—►
Indicates that the syntax diagram is continued to the next line.

►—
Indicates that the syntax is continued from the previous line.

Indicates the end of the syntax diagram.

Syntax items

Syntax diagrams contain many different items. Syntax items include:

- Keywords - a command name or any other literal information.
- Variables - variables are italicized, appear in lowercase, and represent the name of values you can supply.
- Delimiters - delimiters indicate the start or end of keywords, variables, or operators. For example, a left parenthesis is a delimiter.
- Operators - operators include add (+), subtract (-), multiply (*), divide (/), equal (=), and other mathematical operations that may need to be performed.
- Fragment references - a part of a syntax diagram, separated from the diagram to show greater detail.
- Separators - a separator separates keywords, variables or operators. For example, a comma (,) is a separator.

Note: If a syntax diagram shows a character that is not alphanumeric (for example, parentheses, periods, commas, equal signs, a blank space), enter the character as part of the syntax.

Keywords, variables, and operators may be displayed as required, optional, or default. Fragments, separators, and delimiters may be displayed as required or optional.

Item type
Definition

Required
Required items are displayed on the main path of the horizontal line.

Optional
Optional items are displayed below the main path of the horizontal line.

Default
Default items are displayed above the main path of the horizontal line.

Syntax examples

The following table provides syntax examples.

Table 1. Syntax examples

Item	Syntax example
Required item. Required items appear on the main path of the horizontal line. You must specify these items.	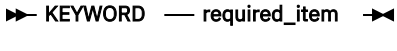
Required choice. A required choice (two or more items) appears in a vertical stack on the main path of the horizontal line. You must choose one of the items in the stack.	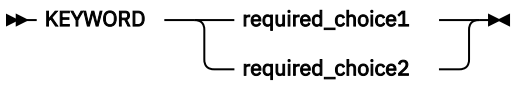
Optional item. Optional items appear below the main path of the horizontal line.	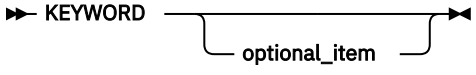
Optional choice. An optional choice (two or more items) appears in a vertical stack below the main path of the horizontal line. You may choose one of the items in the stack.	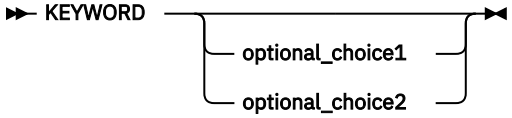

Table 1. Syntax examples (continued)

Item	Syntax example
Default. Default items appear above the main path of the horizontal line. The remaining items (required or optional) appear on (required) or below (optional) the main path of the horizontal line. The following example displays a default with optional items.	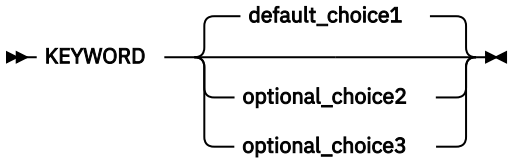
Variable. Variables appear in lowercase italics. They represent names or values.	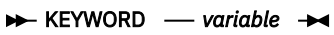
Repeatable item. An arrow returning to the left above the main path of the horizontal line indicates an item that can be repeated. A character within the arrow means you must separate repeated items with that character. An arrow returning to the left above a group of repeatable items indicates that one of the items can be selected,or a single item can be repeated.	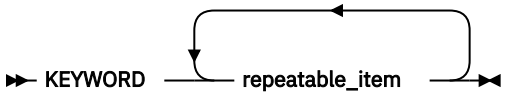 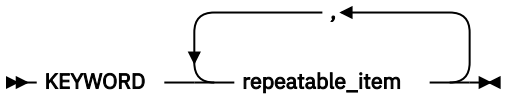
Fragment. The fragment symbol indicates that a labelled group is described below the main syntax diagram. Syntax is occasionally broken into fragments if the inclusion of the fragment would overly complicate the main syntax diagram.	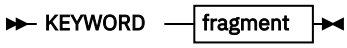 **fragment** 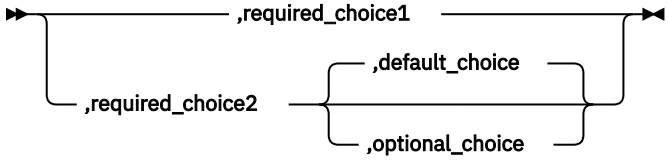

How to send your comments to IBM

We invite you to submit comments about the z/OS product documentation. Your valuable feedback helps to ensure accurate and high-quality information.

Important: If your comment regards a technical question or problem, see instead "If you have a technical problem" on page xv.

Submit your feedback by using the appropriate method for your type of comment or question:

Feedback on z/OS function
If your comment or question is about z/OS itself, submit a request through the IBM RFE Community (www.ibm.com/developerworks/rfe/).

Feedback on IBM Knowledge Center function
If your comment or question is about the IBM Knowledge Center functionality, for example search capabilities or how to arrange the browser view, send a detailed email to IBM Knowledge Center Support at ibmkc@us.ibm.com.

Feedback on the z/OS product documentation and content
If your comment is about the information that is provided in the z/OS product documentation library, send a detailed email to mhvrcfs@us.ibm.com. We welcome any feedback that you have, including comments on the clarity, accuracy, or completeness of the information.

To help us better process your submission, include the following information:

- Your name, company/university/institution name, and email address
- The following deliverable title and order number: z/OS RMF Data Gatherer User's Guide, SC27-4934-40
- The section title of the specific information to which your comment relates
- The text of your comment.

When you send comments to IBM, you grant IBM a nonexclusive authority to use or distribute the comments in any way appropriate without incurring any obligation to you.

IBM or any other organizations use the personal information that you supply to contact you only about the issues that you submit.

If you have a technical problem

If you have a technical problem or question, do not use the feedback methods that are provided for sending documentation comments. Instead, take one or more of the following actions:

- Go to the IBM Support Portal (support.ibm.com).
- Contact your IBM service representative.
- Call IBM technical support.

Summary of changes

This information includes terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations for the current edition are indicated by a vertical line to the left of the change.

Summary of changes for Version 2 Release 4

Notice: This publication contains information that previously appeared in *z/OS RMF User's Guide*, SC34-2664.

This publication presents usage information specific to RMF data gathering. For usage information specific to RMF reporting, see *z/OS RMF Reporter User's Guide*, SC27-4936.

New

- Information about the EADM data gatherer option has been added in “EADM ” on page 77.

Changed

- The following tables are updated to add the CRYPTO option: Table 6 on page 73 and “Default gatherer session options” on page 74.
- The following topics are updated to include the CRYPTO option: Table 7 on page 74 and “Description of Monitor III data gatherer options” on page 75.
- “Modifying the data set support options” on page 87 was updated for the default and maximum values for WHOLD.
- “Starting data set support” on page 86 was updated with the new WHOLD and WSTOR default values.
- “Controlling data set recording” on page 84 was updated for the new WHOLD default value.
- The figure Figure 4 on page 80 was updated for WHOLD and WSTOR.
- The “WSTOR” on page 83 definition was updated.
- The entry for WSTOR was updated in “DATASET” on page 76 within the topic “Description of Monitor III data gatherer options” on page 75.
- “ERBRMF04” on page 25 has been updated with revised information for DATASET(WHOLD) and WSTOR.
- The SCM data gatherer option has been deprecated and is replaced by the EADM option. See “EADM ” on page 77.

Part 1. Introduction

The introduction provides an overview of the RMF portfolio.

- Data Gathering with Monitor I, Monitor II, and Monitor III
- RMF reporting capabilities
- Cross-sysplex data sharing with Sysplex Data Server

Chapter 1. RMF - your performance management tool

Many different activities are required to keep your system running smoothly, and to provide the best service on the basis of the available resources and workload requirements. The operator, the administrator, the system programmer, or the performance analyst will do these tasks. RMF is the tool that helps each of these people do the job effectively.

RMF consists of several components:

- Monitor I - Monitor II - Monitor III
- Postprocessor
- Spreadsheet Reporter
- Sysplex Data Server
- Distributed Data Server for z/OS
- Distributed Data Server for AIX® and Linux systems (RMF XP)
- z/OSMF Performance Monitoring

These components work together in providing the capabilities you need for performance management:

- Gathering data
- Reporting data
- Accessing data across the sysplex

Gathering data

RMF gathers data using three monitors:

- short-term data collection with Monitor III
- snapshot monitoring with Monitor II
- long-term data gathering with Monitor I and Monitor III

The system operator starts all monitors as non-interactive (background) sessions with a variety of options that determine what type of data is collected and where it is stored. The data gathering functions run independently on each system, but each monitor can be started sysplex-wide by one operator command.

You can run data gathering on each z/OS system and use the RMF Sysplex Data Server to have all data available on the one system on which you run your performance management tasks.

Short-term data collection with Monitor III

The Monitor III gatherer session has a typical gathering cycle of one second, and consolidated records are written for a range which is typically set to 100 seconds.

You can collect short-term data and continuously monitor the system status to solve performance problems. You get actual performance data (response times, execution velocity) on a very detailed level for later comparison with performance policy goals.

You can collect data that indicate how fast jobs or groups of jobs are running — this is called **workflow** or **speed**. You also get data that show how resource-intensive jobs are using the processor, the DASD devices, and the storage — the reports describe this under the term **using**.

There is also information about delays, which are important indicators of performance problems. This simplifies comparison of reports created from Monitor I and Monitor III data.

Snapshot monitoring with Monitor II

The scope of Monitor II data gathering is mainly related to single address spaces or resources, giving snapshots of the current status. You can collect data about address space activities and resource consumption, and about processor, DASD volume, and storage activities and utilization.

With Monitor II, it is also possible to monitor one specific job or volume continuously.

Long-term data gathering with Monitor I and Monitor III

Monitor I and Monitor III provide long-term data collection about system workload and resource utilization, and cover all hardware and software components of your system: processor, I/O device and storage activities and utilization, as well as resource consumption, activity and performance of groups of address spaces.

Data is gathered for a specific cycle time, and consolidated data records are written at a specific interval time. The default value for data gathering is one second and for data recording 30 minutes. You can select these options according to your requirements and change them whenever the need arises.

The SMF synchronization function ensures that records are written from all monitors in the sysplex for the same intervals.

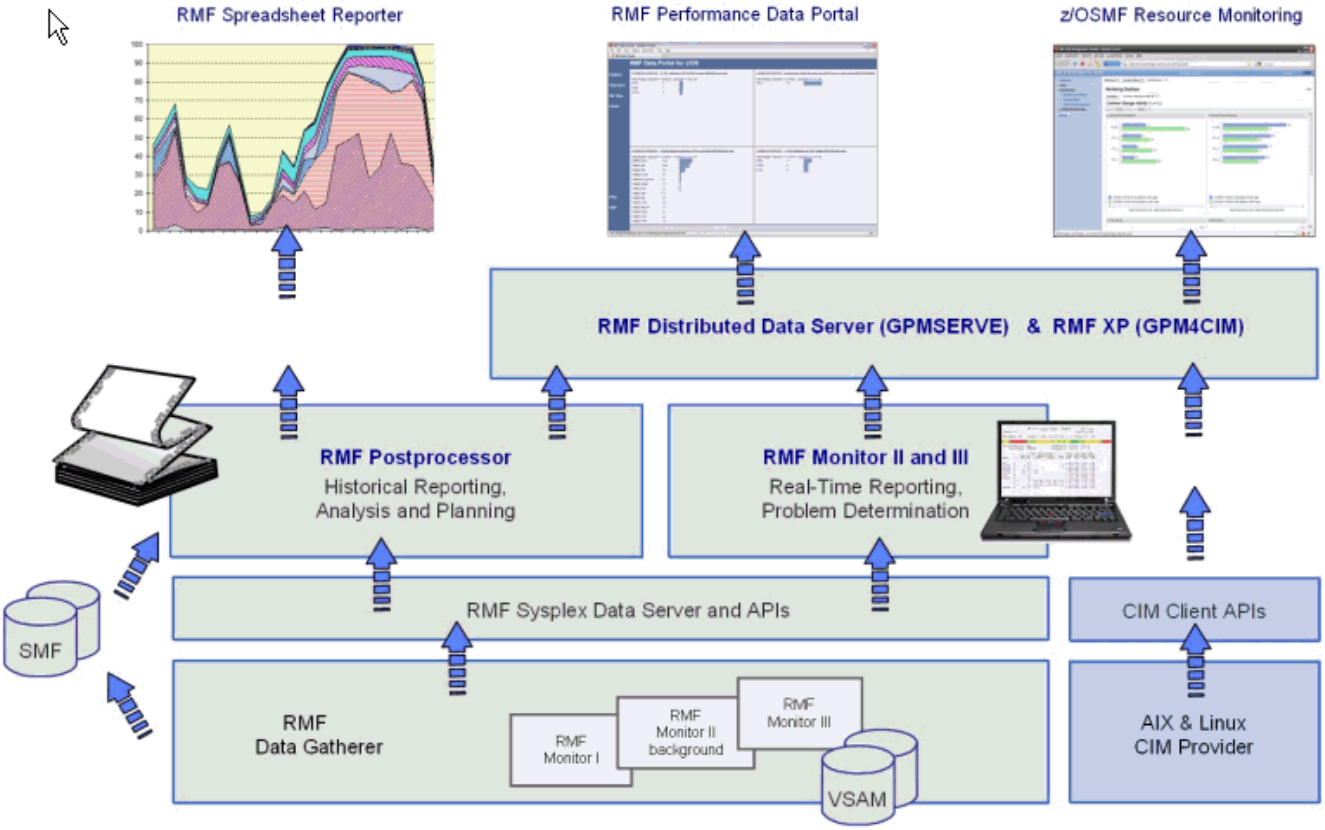

Figure 1. RMF - Your Performance Management Tool

Monitored activities

RMF stores data in two types of record:

- All three monitors write SMF records (type 70 — type 79) if you define the appropriate SMF recording options.
- In addition, Monitor III writes VSAM records to in-storage buffers or into RMF-owned VSAM data sets.

Table 2 on page 5 provides the following information:

- Displays the SMF record types of all records that are written by gatherer sessions

- Indicates all Monitor III data stored in VSAM data sets
- Shows all corresponding reporting components

Table 2. Monitored activities and SMF record types

Gathering				Activity	Reporting			
Short-term Mon III		Snapshot Mon II	Long-term Mon I		Interactive Mon III	Snapshot Mon II	Real-time Mon I	Long-term Post-processor
SMF	VSAM	SMF	SMF					
	*	79.1/2/5		Address space	*	*		*
	*		74.5	Cache	*			*
	*	79.12	73	Channel path	*	*	*	*
74.4	*			Coupling facility	*			*
	*		70.2	Cryptographic hardware	*		*	*
	*	79.9	74.1	Device	*	*	*	*
74.10	*			EADM Activity	*			*
	*			Enclave	*			
	*	79.7	77	Enqueue	*	*	*	*
			74.8	Enterprise Storage Server (ESS)				*
			74.7	FICON director				*
		79.15		IRLM long locks		*		
	*	79.14	78.3	I/O queuing	*	*	*	*
		79.11	75	Page data set		*	*	*
		79.4	71	Paging		*	*	*
74.9	*			PCIE Activity	*			*
	*	79.3	70.1	Processor	*	*	*	*
		79.6		Reserve		*		*
72.5				Serialization Delay				*
72.4	*	79.3		Storage	*	*		*
			76	System counters			*	*
74.3/6	*			UNIX	*	*		*
	*		78.2	Virtual storage	*		*	*
	*		72.3	Workload Service classes and report classes	*			*
74.2	*			XCF	*			*
	*			zFS	*			

Reporting data

RMF provides the following reporting capabilities:

- Short-term interactive performance analysis with Monitor III
- Snapshot reporting with Monitor II
- Long-term overview reporting with the Postprocessor

- Viewing Postprocessor reports on spreadsheets with Spreadsheet Reporter
- Cross-sysplex performance monitoring with IBM z/OS Management Facility (z/OSMF)
- Sysplex-wide data providing with RMF Distributed Data Server (DDS)
- Resource monitoring of systems running AIX or Linux with RMF Cross Platform (RMF XP)

For more information about the particular RMF reporting components, see *z/OS RMF Reporter User's Guide*.

Accessing data across the sysplex

Read the following subtopics to learn how to access performance data across the sysplex:

- "RMF Sysplex Data Server" on page 6
- "Sysplex data services for SMF data" on page 7
- "Sysplex data service for Monitor III data" on page 7
- "Sysplex data gathering service for Monitor II data" on page 7

RMF Sysplex Data Server

The RMF Sysplex Data Server is a distributed RMF function. It is started as an identical copy on each system of the sysplex. Each copy of the data server communicates with all other copies in the sysplex. RMF uses this sysplex communication method to provide access to distributed RMF measurement data from any point in the sysplex.

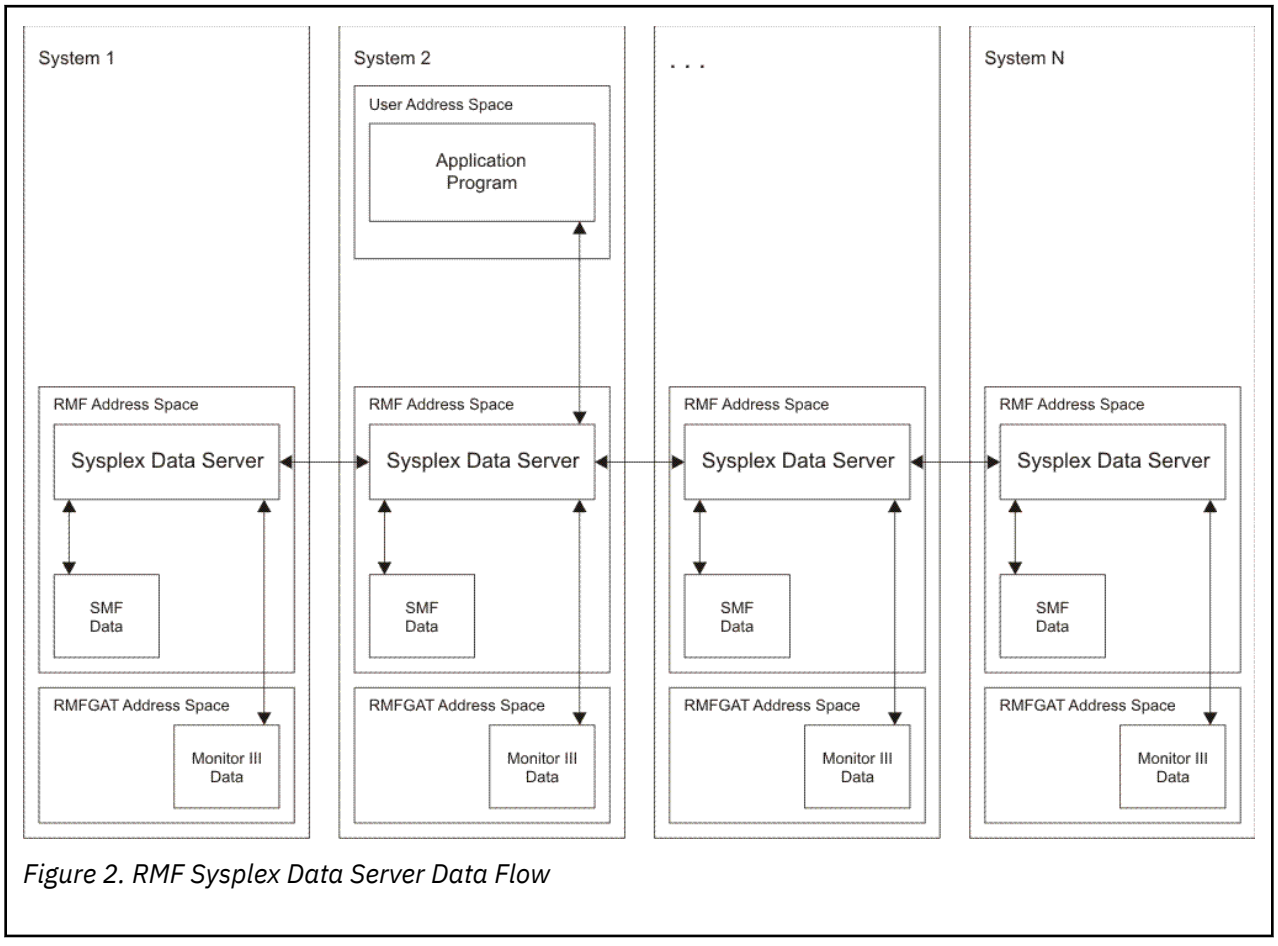

Figure 2. RMF Sysplex Data Server Data Flow

The RMF Sysplex Data Server is always active when the RMF address space is running.

You can access all types of RMF and SMF data collected in the sysplex by using RMF Sysplex Data Server programming interface services. These are invoked as callable services by the RMF reporter sessions themselves or other applications, and can access:

- Monitor I, II and III SMF data
- Monitor III VSAM data
- SMF data of any other type

To call the RMF services for SMF data, you need authorization to access the SMF data.[1] For details, please see “Controlling access to RMF data for the sysplex data services” on page 13.

Sysplex data services for SMF data

RMF or other products store SMF data in a wrap-around buffer. You can choose to create a RMF Sysplex Data Server's SMF buffer when you start RMF. The size of the buffer and the types of SMF records stored in it can be specified as an RMF startup parameter. The RMF sysplex data services return SMF data when the RMF Sysplex Data Server's SMF buffer exists on at least one system in the sysplex, which need not be the system on which the calling program is running. The Data Server returns data only from systems in which data buffers have been created.

Sysplex data service for Monitor III data

You can access data collected by Monitor III data gatherer sessions using the RMF Monitor III Sysplex Data Retrieval Service. Any application program can specify the name of the system from which the Monitor III data is requested. Analogous to SMF data, Monitor III data can be returned from those systems where the Monitor III data gatherer session is active.

Sysplex data gathering service for Monitor II data

Your application program can use this service to create and retrieve Monitor II SMF records (type 79). You need not have a Monitor II background session running on the system from which you request the data. Note the difference between this service and the data service for SMF data, which collects only records created by active monitor sessions.

[1] Authorization of application programs is provided by the z/OS Security Server RACF®, or products with similar functions, that define the user group authorized to access measurement data returned by the RMF Sysplex Data Server callable services.

The services may be invoked by programs running under any PSW key and in problem state, like the Postprocessor and Monitor III reporter sessions.

Part 2. Administration

Administration is what you have to do after installing RMF and before you start using it for measuring resources. The administrator creates the prerequisites that the daily user takes for granted, like setting up job control procedures and defining standard data sets for the installation.

Unlike installation, administration is typically an on-going task, though not as frequent as resource measurement. Start with administrative effort after installation, and continue as the needs of the users change with changing conditions in the system.

An RMF administrator can:

- Define **system parameters** and **access definitions** being required for smoothly running gathering functions.
- Update the **RMF cataloged procedure** to define the gatherer options, and to set default values for the SMF wrap-around data buffer, in which RMF monitors store the data they collect.
- Synchronize **SMF recording intervals** with data collection intervals of the RMF monitors, to obtain comparable measurements.
- Define **VSAM data sets** for storing data from Monitor III gatherer sessions.
- Define the **parmlib members** for the RMF monitors. These determine the default options for the respective monitors, so their contents should be agreed upon between administrator and performance analyst. A parmlib member for each monitor is provided with RMF, but can be modified as required. The options that can be included in the members are described in detail in Part 4, “Data gathering reference,” on page 47.

Most of these tasks have to be performed only once during the initial customization of RMF according to the requirements of your installation. An ongoing task is the migration from one release to the next one. Therefore, these migration steps are described separately to highlight the differences between releases and the actions which might be required to ensure that you exploit the new functions that RMF is offering with each new release.

Chapter 2. Setting up RMF

After installing RMF, you must perform the following administration tasks:

- The tasks for migrating to the release described in this document
- The steps for activating RMF functions
- The JCL procedure for starting the RMF control session
- The JCL procedure for starting the Monitor III gatherer session
- The definition of Monitor III gatherer VSAM data sets
- Synchronization with SMF data recording
- The parmlib members that contain your system's standard gatherer options

Migrating from previous releases

If you have installed and activated RMF in a previous release, in most cases it is not required to change anything in the procedures and parameters you are using. Typically, new gathering options will be activated automatically, and special considerations might be necessary only if you are using customized parmlib members for data gathering.

All required or optional migration actions for RMF are documented in *z/OS Upgrade Workflow*.

Customizing the system environment

This section describes the required tasks for customizing the system environment in order to ensure a proper functioning of RMF on this system.

Define RMF library authorization

All RMF load modules reside in the two libraries SYS1.SERBLINK and SYS1.SERBLPA. If you are activating RMF for the first time, you have to define these libraries as APF authorized libraries. You can choose to do it either with or without an IPL.

To activate RMF with an IPL:

1. Add the SERBLINK library to the link list
2. Add the SERBLINK library to the APF list
3. Add the SERBLPA library to the LPA list
4. IPL the system

To activate RMF without an IPL:

1. Add the SERBLINK library to a dynamic link list
2. Change the APF format to dynamic, if it is not already dynamic
3. Add the SERBLINK library to the dynamic APF list
4. Add the SERBLPA library to Dynamic LPA
5. Issue SETPROG commands to make the changes effective

For more information about adding libraries to the link, APF, and LPA lists with or without an IPL, see *z/OS MVS Initialization and Tuning Reference*. For information about the syntax of the SETPROG command, see *z/OS MVS System Commands*.

Check the program properties table (PPT)

z/OS provides two default entries in the PPT for the RMF modules ERBMFMFC and ERB3GMFC. You should run with the defaults provided in the PPT, or the results will be unpredictable. The default entries include:

- Non-swappable
- System task
- No protection key
- No processor affinity

Any user modifications to those entries require to specify a PPT entry for ERBMFMFC and ERB3GMFC in a SCHEDxx parmlib member, which must include the RMF defaults and user overrides. Here is an SCHEDxx example:

```
/* PPT Entry for RMF (RMF Control/Monitor I)            */
PPT PGMNAME(ERBMFMFC)          /*PROGRAM NAME           */
    CANCEL                     /*CAN BE CANCELLED       */
    NOSWAP                     /*NON-SWAPPABLE          */
    NODSI                      /*NO DATA SET INTEGRITY  */
    PASS                       /*NO PASSWORD BYPASS     */
    SYST                       /*SYSTEM TASK, NOT TIMED */
    AFF(NONE)                  /*NO PROCESSOR AFFINITY  */
/* PPT Entry for RMFGAT (Monitor III data gatherer)     */
PPT PGMNAME(ERB3GMFC)          /*PROGRAM NAME           */
    CANCEL                     /*CAN BE CANCELLED       */
    NOSWAP                     /*NON-SWAPPABLE          */
    NODSI                      /*NO DATA SET INTEGRITY  */
    PASS                       /*NO PASSWORD BYPASS     */
    SYST                       /*SYSTEM TASK, NOT TIMED */
    AFF(NONE)                  /*NO PROCESSOR AFFINITY  */
```

Note: Do **not** specify a protection key for these entries.

Remove ERBMFRES

If you are installing RMF on a system that already has ERBMFRES (Memory Termination Resource) in the resource manager list, you should remove it or you will experience performance degradation.

This resource manager list (table IEAVTRML) is located in the load module IGC0001C.

Global performance data control authority

This control limits the ability of a logical partition to retrieve global performance data for other logical partitions. RMF can report on CPU utilization data and Input/Output Processor (IOP) data for all logical partitions in the configuration only if this security option is selected. If not selected, RMF reports only CPU utilization data for its own logical partition. In addition, gathering of channel measurements requires control authority.

This option is selected per default in the logical partition security controls on the Hardware Management Console. For more information refer to the appropriate *PR/SM Planning Guide*.

Specifying access definitions

Read the following subtopics for information on how to grant the required access rights for setting up RMF:

- "Define RMF user IDs and ensure access to z/OS UNIX System Services" on page 13
- "Assign started task procedures to user IDs" on page 13
- "Controlling access to RMF data for the sysplex data services" on page 13

Define RMF user IDs and ensure access to z/OS UNIX System Services

At first, you should define user IDs that are associated with RMF data gathering. We recommend to define two user IDs that relate to the two started tasks that RMF provides for data gathering. (Of course, you may decide to define only one user ID that is assigned to both started tasks).

Because RMF started tasks use UNIX System Services or resources, the procedures must be defined to the security program. For example, the Monitor III gatherer (RMFGAT) must be defined so that it can obtain the correct data and can use the required UNIX system services.

The following example contains RACF commands to define the two RMF user IDs, to give them an OMVS user ID (UID), and to designate the root directory as its home directory:

```
ALG omvsgrp OMVS(GID(2))
ADDUSER RMF      DFLTGRP(omvsgrp) OMVS(UID(nnn) HOME('/'))
ADDUSER RMFGAT   DFLTGRP(omvsgrp) OMVS(UID(nnn) HOME('/'))
```

In the above example, nnn can be any number, but must not be zero. For details, see *z/OS UNIX System Services Planning*.

Assign started task procedures to user IDs

RMF provides four different started tasks of which two are associated with data gathering. The remaining two tasks are related to the Distributed Data Server and are discussed in *z/OS RMF Reporter User's Guide*. In this step, you define the two data-gathering started tasks to RACF and assign them to the RMF user IDs that you created in the previous step.

```
RDEFINE STARTED RMF.*      STDATA(USER(RMF)      TRUSTED(YES))
RDEFINE STARTED RMFGAT.*   STDATA(USER(RMFGAT)   TRUSTED(YES))
SETROPTS RACLIST(STARTED) REFRESH
```

Controlling access to RMF data for the sysplex data services

Users of applications that call sysplex data services to access data from the RMF Sysplex Data Server's SMF buffer must have RACF authorization.

RMF has defined a RACF resource profile of class FACILITY called ERBSDS.SMFDATA to control access to SMF data in the RMF Sysplex Data Server's SMF buffers. Every user accessing the SMF records in this SMF buffer must be authorized.

ERBSDS.SMFDATA
Controls access to SMF data in the SMF buffer by the ERBDSQRY service (Query Available Sysplex SMF Data) or the ERBDSREC service (Request Sysplex SMF Record Data).

RMF does not perform mandatory access checks for Monitor II data (accessed by the ERB2XDGS service) and Monitor III set-of-samples data (accessed by the ERB3XDRS service). If you want to protect this data, define RACF resource profiles called ERBSDS.MON2DATA and ERBSDS.MON3DATA in the FACILITY class. If you do not define a profile, RACF does not restrict any user ID from invoking the mentioned sysplex data services:

ERBSDS.MON2DATA
Controls access to Monitor II SMF type 79 data by the ERB2XDGS and ERBSMFI services.

ERBSDS.MON3DATA
Controls access to Monitor III set-of-samples data by the ERB3XDRS service.

If the same group of users takes advantage of all RMF sysplex data services, you can work with the generic profile ERBSDS.*.

Controlling the invocation of data reduction exit routines

RMF controls the names of the data reduction exit routines that are provided by callers of RMF Monitor II Sysplex Data Gathering service ERB2XDGS or Monitor III Sysplex Data Retrieval service ERB3XDRS. Invocation of these exit routines is controlled in the following ways:

- If an ***authorized*** caller is running in supervisor state, in system state, or APF authorized, then it can use ***trusted*** exit names.
- If an ***unauthorized*** caller specifies ***untrusted*** but ***approved*** exit names, then you can specify the following access controls:
  1. Define the RACF resource profile ERBSDS.MON2EXIT.<exit_name> to the class FACILITY. The use of the data reduction exit routine with the name specified with <exit_name> will be restricted to those ERB2XDGS callers who have been authorized to this RACF resource profile.
  2. Define the RACF resource profile ERBSDS.MON3EXIT.<exit_name> to the class FACILITY. The use of the data reduction exit routine with the name specified with <exit_name> will be restricted to those ERB3XDRS callers who have been authorized to this RACF resource profile.
- If an ***unauthorized*** caller specifies ***untrusted*** exit names that are ***not approved***, then the ERB2XDGS service as well as the ERB3XDRS service provide return code 16 and reason codes 86 or 87, and RACF issues message ICH408I indicating the exit name that caused the security violation.

For more information about the involved RMF sysplex data services and the description of their return an reason codes, refer to *z/OS RMF Data Gatherer Programmer's Guide*.

Security server example

This information unit presents a security server coding example using RACF to achieve the required access and invocation control as described in "Controlling access to RMF data for the sysplex data services" on page 13 and "Controlling the invocation of data reduction exit routines" on page 13.

1. To activate the resource class FACILITY:

   ```
   SETROPTS CLASSACT(FACILITY) GENCMD(FACILITY) GENERIC(FACILITY)
   ```

2. To define the profile:

   ```
   RDEFINE FACILITY <profile> UACC(NONE)
   ```

   where <profile> is one of the following profile names:

   - ERBSDS.SMFDATA (mandatory)
   - ERBSDS.MON2DATA (optional)
   - ERBSDS.MON3DATA (optional)
   - ERBSDS.MON2EXIT.<exit_name> (mandatory for untrusted exit names)
   - ERBSDS.MON3EXIT.<exit_name> (mandatory for untrusted exit names)
   - or the generic profile name ERBSDS.*.

   The name specified with <exit_name> denotes the data reduction exit routine used with the corresponding sysplex service.

3. To grant the user ID of the application program READ access:

   ```
   PERMIT <profile> CLASS(FACILITY) ID(<userid>) ACC(READ)
   ```

4. Activate changes:

   ```
   SETROPTS  RACLIST(FACILITY) REFRESH
   ```

Checklist for access to sysplex data services

If you want to prevent unauthorized access to the sysplex data services, you can use the following checklist to ensure that you completed all required tasks:

- define the profiles ERBSDS.SMFDATA, ERBSDS.MON2DATA and ERBSDS.MON3DATA to the FACILITY class to protect access to the related sysplex data services
- or work with the generic profile ERBSDS.* and have generic profile checking active

Setting up the RMF control session including Monitor I and Monitor II

You should perform the following steps to ensure correct data gathering with Monitor I and Monitor II:

- "Customizing the RMF control session" on page 15
- "Specifying priority for RMF" on page 16
- "Storing gatherer defaults" on page 16

Customizing the RMF control session

IBM provides the cataloged RMF procedure which is necessary to start RMF. The procedure is stored in SYS1.PROCLIB(RMF), and you can modify it according to your requirements.

The RMF control session is the base for data gathering through the different monitors, especially for Monitor I and Monitor II. If you want to gather data with Monitor III, you need in addition procedure RMFGAT (see "Setting up the Monitor III gatherer session RMFGAT" on page 16).

This example shows the RMF procedure as supplied:

```
//IEFPROC   EXEC  PGM=ERBMFMFC,REGION=256M,
//          PARM=''
```

PARM
can be used for:

- specifying the SMF buffer options to be used by the RMF Sysplex Data Server. The format of this option is described in "Controlling the SMF buffer" on page 31. The defaults mean that specifying PARM='SMFBUF' is equivalent to:

  ```
  PARM='SMFBUF(SPACE(32M),RECTYPE(70:78))'
  ```

  You can override the values specified or defaulted here by using the **SMFBUF** option on the **START RMF** command when starting RMF.
- providing automatic sysplex-wide management for the Distributed Data Server (PARM='DDS', see Starting the Distributed Data Server in *z/OS RMF Reporter User's Guide*)
- specifying the Monitor I gatherer session options, for example, PARM='MEMBER(10),NOCACHE', (see Chapter 6, "Long-term data gathering with Monitor I," on page 49)

RMF reads its ERBRMFxx members from the parmlib concatenation as defined in the LOADnn parmlib member, and then frees the data set in which they were found.

To have RMF read the ERBRMFxx members from a specific single data set, use a cataloged procedure in the following form:

```
//IEFPROC   EXEC  PGM=ERBMFMFC,REGION=256M,
//          PARM=''
//IEFPARM   DD    DSN=parmlibname,DISP=SHR
```

IEFPARM
Identifies the data set containing session options. If you specify an IEFPARM DD statement in the procedure, RMF does not use the logical parmlib concatenation.

To start the SMF data buffer on each system in your sysplex, store the procedure in the common proclib as follows:

```
//IEFPROC   EXEC  PGM=ERBMFMFC,REGION=256M,
//          PARM='SMFBUF'
```

Specifying priority for RMF

The started tasks RMF and RMFGAT must have the second-highest priority in the system, next to the system address spaces. Use the WLM application to put RMF and RMFGAT in service class SYSSTC to ensure that its dispatching priority will always be above any installation-defined service class. If the priority is too low, it can happen that RMF is not dispatched when its interval time expires, with the consequence that data collection for jobs running with higher priority is incomplete, or that any event processing cannot be performed. This could result either in incorrect measurement reports, or in common storage shortages, which might lead to an IPL.

Storing gatherer defaults

The Monitor I and Monitor II gatherer sessions require several parameters to define the type of data to be gathered. These parameters are stored in parmlib members, and are used when you start the gatherer session. The parmlib members supplied with RMF contain meaningful values, but you can change these to suit your purposes, or you can create new parmlib members and have them used at session start.

The parmlib members with the supplied defaults are described in “Storing gatherer options” on page 19.

Setting up the Monitor III gatherer session RMFGAT

Preparation of data gathering with Monitor III requires the following steps:

- “Defining VSAM data sets” on page 16
- “Ensuring common storage tracking” on page 17

IBM provides the cataloged procedure needed to start the Monitor III gatherer session. It is stored in SYS1.PROCLIB(RMFGAT):

```
//IEFPROC  EXEC PGM=ERB3GMFC,REGION=256M,TIME=1440
```

RMF dynamically allocates the Monitor III gatherer message file RMFM3III to SYSOUT=A, but you can insert a DD statement in the RMFGAT procedure to preallocate it.

Since RMF is running with NODSI setting in its PPT entry, be aware that RMF does not hold any data set ENQUEUE (major name=SYSDSN, minor name=dsname) for the data sets allocated by the RMF procedure. A missing ENQUEUE can mislead functions like HSM, that rely on ENQUEUE for data set protection.

In a system without an active JES2 or JES3, you **must** make this preallocation before you start the Monitor III data gatherer (see “Starting RMF without JES” on page 30). You can use a DD DUMMY statement if you do not wish to store the gatherer messages.

Beginning with z/OS V2R2, the RMF Monitor III gatherer uses high virtual private storage. Ensure that the MEMLIMIT for the RMFGAT started task is at least 2G. If your active SMFPRM*xx* parmlib member specifies a MEMLIMIT less than 2G (the default is 2G), add MEMLIMIT=2G to your RMFGAT started task procedure.

Defining VSAM data sets

The Monitor III data gatherer writes records (sets of samples) to a large storage buffer, or optionally, to user-defined VSAM data sets. Without the VSAM data sets, the data is overwritten as soon as the buffer is filled. If you define VSAM data sets, you can save large amounts of information, and you can reuse the VSAM data sets as RMF continuously records data over time.

You can define up to 100 data sets for use with the data gatherer. You should define at least two data sets, because the gatherer deletes all data in a data set before writing to it, so a single data set would be emptied immediately after it was filled. RMF can keep about 1100 sets of samples in one data set. Based on practical experience, we recommend to define six VSAM data sets, each with 50 cylinders disk space. On small and medium systems, this will allow for about two days of data.

Note: If you need to change the amount of space later to meet your installation's needs, we recommend adding more data sets, but not making the individual data sets larger. Increasing the size of the data sets may cause RMF to run out of index entries and be unable to fill the additional space.

Have a look at the Data Index (see Using the Data Index (DI) in *z/OS RMF Reporter User's Guide*) which tells you exactly the time range of the data that is available in your VSAM data sets. This can help you in defining the appropriate number of data sets.

Sysplex considerations

If you run RMF in a sysplex, it is recommended to select names for the VSAM data sets with the MVS™ system name being part of the data set name. Then you can easily use the capability of symbolic names to specify your parmlib members. Please refer to "Generalizing parmlib members" on page 20 for details.

Defining VSAM clusters

You must define the VSAM data sets to be used for recording data before you start a Monitor III data gatherer session. When you specify a data set on the DATASET option, you must use the `dsname` you define on the `NAME` parameter of the `DEFINE CLUSTER` statement.

You can use the CLIST ERBVSDEF, shipped in SYS1.SERBCLS, to define the data sets.

```
ERBVSDEF vsam_dsn VSAMVOL(volume) [TRACKS(num_tracks)]
```

where:

vsam_dsn
is the name of the Monitor III VSAM data set to be allocated.

volume
is the volume on which the VSAM data set is to be allocated, this parameter is required for systems on which SMS is not active.

num_tracks
is the primary extent of the VSAM data set (the default is 150 tracks).

To define a VSAM data set named RMF.MONIII.DS1 on the volume DATA01, enter:

```
ERBVSDEF 'RMF.MONIII.DS1' VSAMVOL(DATA01)
```

"Controlling data set recording" on page 84 tells you how to specify which data sets are to be used for a particular Monitor III gatherer session.

Ensuring common storage tracking

To ensure that the Common Storage report (STORC) provides complete data, it is required that VSM common storage tracking is active. This can be achieved by issuing the command:

```
SET DIAG=01
```

The defaults in the parmlib member DIAG01 are:

```
VSM TRACK CSA(ON) SQA(ON)
```

If VSM common storage tracking is not active, one of the messages ERB617I, ERB618I, or ERB619I will indicate that the report can be incomplete for some jobs.

Synchronizing SMF recording intervals

All RMF monitors write SMF records if you specify the appropriate gatherer options. The Postprocessor can later process these records to create comprehensive reports of either single-system or sysplex

scope. For sysplex reports, the Postprocessor requires all records written by RMF to be synchronized, and for single-system reports, synchronization is recommended. Therefore, you should perform these tasks:

- "Defining SMF record writing" on page 18
- "Defining SMF synchronization" on page 18

Defining SMF record writing

You can specify by SMF options (defined in the SMFPRMxx parmlib member) and Monitor I and Monitor II gatherer options (defined, for example, in ERBRMFxx parmlib members) whether you want to write SMF records during your gathering sessions.

SMF Recording

SMF
Option TYPE of the SYS command in the active SMFPRMxx parmlib member specifies the SMF record types and subtypes that SMF is to collect:

- Monitor I and Monitor III write record types 70 — 78.
- Monitor II writes record type 79.

Example:

```
SYS(TYPE(...,72,...))         /* write SMF record type 72  */
```

Monitor I / Monitor II
Option RECORD in ERBRMFxx parmlib member specifies SMF record collection.

Monitor III
Automatic record writing if enabled via the SMF option TYPE (in SMFPRMxx).

SMF provides specific user exits to control data collection. Ensure that you do not suppress the writing of RMF records if you want to create Postprocessor or other reports.

See *z/OS MVS Initialization and Tuning Reference* for details.

Suppressing SMF record writing

If RMF per default writes an SMF record type or subtype that you do not want to be written, you can use one of the following methods to suppress SMF recording:

- Use the SUBSYS command in the SMFPRMxx parmlib member. The SUBSYS specification overrides the SYS specification. For example, if you have defined `SYS(TYPE(...,72,...))` in your SMFPRMxx parmlib member, you can use `SUBSYS(STC, NOTYPE(72(5)))` to make exceptions to your SYS specification and just exclude gathering of SMF record 72.5 for started tasks like RMF.

  For more information, see *z/OS MVS Initialization and Tuning Reference*.
- Use the system command SETSMF. For more information, see *z/OS MVS System Commands*.

Defining SMF synchronization

SMF provides options that you can use for synchronization of record writing in the sysplex. The SMF options are:

INTVAL(mm)
SMF global recording interval - default is 30 minutes

SYNCVAL(mm)
Synchronization with the hour - default is 00

If you use the default values, this means that SMF records will be written every 30 minutes at the full and the half hour.

Monitor I has these options that specify when to write SMF records:

SYNC(SMF)
Synchronization with SMF - this is the default and means that records will be written as specified with INTVAL and SYNCVAL options.

SYNC(RMF,mm)
RMF synchronization with the hour

NOSYNC
No synchronization

INTERVAL(mm)
Interval length - this value is ignored with SYNC(SMF)

The synchronization of SMF records written by Monitor III is defined by the SMF and Monitor I options:

Monitor I active
Monitor III has the same synchronization as Monitor I

Monitor I inactive
Monitor III has the global SMF synchronization (defined by INTVAL and SYNCVAL)

Note:

1. If you intend to create Postprocessor sysplex reports, you must use the same SYNC values on all systems in the sysplex. Do not use NOSYNC on any of the systems, in this case.
2. Nevertheless, different interval lengths are acceptable (but not recommended). The Postprocessor will use the smallest common multiplier to determine the interval length.

   For example, if you have intervals of 10 minutes on SYSA and 15 minutes on SYSB, a sysplex report will be generated every 30 minutes (taking three intervals from SYSA and two intervals from SYSB).

Recommendation

Use the following values to synchronize SMF record writing:

SMF
INTVAL(nn) SYNCVAL(00) where nn can be 05, 10, 12, 15, 20, 30 or 60

Monitor I
SYNC(SMF)

For information about SMF record format and printing, see *z/OS RMF Data Gatherer Programmer's Guide*. See *z/OS MVS System Management Facilities (SMF)* for descriptions and formulas of the fields for each SMF record that RMF produces.

Storing gatherer options

Perform the following tasks to tailor RMF data gathering according to your requirements:

- "Naming parmlib option members" on page 20
- "Generalizing parmlib members" on page 20
- "Defining parameters for Monitor I" on page 22
- "Defining parameters for Monitor II" on page 24
- "Defining parameters for Monitor III" on page 25

You can choose the options for each gatherer session in three ways:

- By accepting the RMF defaults
- By specifying options on a system command
- By storing a list of session options in a parmlib member

This chapter tells you how to specify session options in a parmlib member.

Naming parmlib option members

The parmlib members containing gatherer session options must be named ERBRMFxx, where xx is two alphameric characters. Each data gatherer has a MEMBER option, which allows you to specify the parmlib member from which the options are to be taken for the current session. For example, specifying MEMBER(08) causes RMF to use the options in the ERBRMF08 parmlib member.

If you do not specify a MEMBER option, RMF uses a particular default parmlib member for each type of gatherer session:

- ERBRMF00 for Monitor I
- ERBRMF01 for Monitor II
- ERBRMF04 for Monitor III

These members are supplied with RMF, as are two alternative members:

- ERBRMF02 for Monitor I
- ERBRMF03 for Monitor II

You can use the default and alternative members as they are, or you can alter them to suit your needs. You can also create new parmlib members from scratch, following the naming convention of ERBRMFxx. For the options and their syntax, see Chapter 6, “Long-term data gathering with Monitor I,” on page 49, Details of report commands in *z/OS RMF Reporter User's Guide*, and Chapter 8, “Short-term data gathering with Monitor III,” on page 73, respectively.

Remember that to use any parmlib members other than the defaults, you must specify them on the MEMBER option when starting the respective monitor.

Syntax rules for ERBRMFxx

The following syntax rules apply to the ERBRMFxx parmlib members:

- Use columns 1 through 71. Columns 72 through 80 are ignored.
- Comments begin with /* and end with */.

Parmlib concatenation

With the support of parmlib concatenation in z/OS, it is recommended to define one or more *customer* parmlibs that can be specified in the LOADnn parmlib member. Then you can distinguish between system-supplied members (for example through the SMP/E installation process) which will be stored by default in SYS1.PARMLIB, and customer-modified members in an additional parmlib data set.

If you modify members ERBRMF00 - ERBRMF04 according to your requirements, you should store them in a separate parmlib to avoid that they will be overwritten unintentionally during the installation of an APAR or a follow-on release.

The total number of lines that can be specified in all concatenated ERBRMFxx parmlib members is limited to 448.

Generalizing parmlib members

In a sysplex environment, each individual system has its own parmlib with the corresponding RMF parmlib members. It is often convenient to generate a new system in the sysplex by cloning an existing one, but any references to the system name in, for example, parmlib members, must be altered accordingly.

To make this adaptation automatic, RMF uses the capability of working with symbolic names. They can be defined by you as the user, and there are a number of predefined symbolic names that you can use without further preparation.

The predefined symbolic names &SYSNAME and &SYSCLONE are the most useful for the RMF user. &SYSNAME resolves to the 8-character MVS system name, and &SYSCLONE to the last two non-blank characters of the system name.

RMF supports the use of symbolic names in:

- All RMF parmlib members

  Now, you can use the same parmlib member on each system, if you use symbolic names for system-specific options, as shown in the following examples.

  Example of using symbolic names for system-specific options:

  To ensure that RMF uses different VSAM data sets on each system in the sysplex without the need for different parmlib members, include in the Monitor III parmlib member:

```
⋮
DATASET(START)
DATASET(ADD(SYS1.ERB.&SYSNAME..VSAM1))
DATASET(ADD(SYS1.ERB.&SYSNAME..VSAM2))
⋮
```

  Another example of using symbolic names:

  Assume you have a CICS® address space running on each of your systems in the sysplex and for easy naming you named these address spaces CICS1 (running on system PRD1), CICS2 (on PRD2) and CICS3 (on PRD3).

  If you want to monitor these address spaces with Monitor II in the background, you can specify in your Monitor II parmlib member:

```
⋮
ASRMJ(CICS&SYSCLONE(2:1))
ASDJ(CICS&SYSCLONE(2:1))
ARDJ(CICS&SYSCLONE(2:1))
⋮
```

  Example that shows how to set up gathering options:

  In an environment where several systems have access to one and the same storage subsystem, it is sufficient that the cache data gatherer is started just on one system. Running the gatherer on more than one system creates several copies of identical SMF records type 74-5 (Monitor I) or VSAM records (Monitor III).

  Since RMF has no sysplex control over the gatherer options, it cannot automatically deselect cache gathering on all but one system. To take advantage of shared parmlibs in a sysplex environment, help yourself using the symbolics approach offered by z/OS.

  – Specify an IEASYMxx parmlib member in your LOADxx-member:

```
IEASYM CA
```

  – Define a symbol &CACHEOPT in parmlib member IEASYMCA (assuming that the sysplex is built from z/OS systems running in LPAR partitions):

```
SYSDEF SYMDEF(&CACHEOPT='NOCACHE') /* Global value        */
SYSDEF LPARNAME(PROD1)
       SYMDEF(&CACHEOPT='CACHE')  /* Local value for SYS1 */
```

  – Create a shared RMF parmlib member ERBRMFxx:

```
  ...                              /* any global RMF parms */
  &CACHEOPT.                       /* CACHE or NOCACHE     */
  ...                              /* any global RMF parms */
```

  – Start RMF on all systems using the member option:

```
RO *ALL,S RMF.A,,,(MEMBER(xx))
```

  With this definition, the symbol &CACHEOPT is defined as 'NOCACHE', while on system SYS1, the symbol is resolved as 'CACHE'.

  For details about defining your own symbols, refer to *z/OS MVS Initialization and Tuning Reference*.

- The reply to message ERB306D REPLY WITH OPTIONS OR GO

  You can use symbolic names in the option strings that you type in at the terminal, using the same conventions as in the parmlib members
- The RMF MODIFY command. Again, the options can contain symbolic names, as in the parmlib members. The command is converted automatically during MVS command processing. The system responds to a MODIFY command that contains symbolic names as shown in the following example.

  Example of command with symbolic name, and system response:

```
f rmf,f iii,dataset(add(SYS1.&SYSNAME..DATA))
IEE295I COMMAND CHANGED BY SYMBOLIC SUBSTITUTION
ORIGINAL: F RMF,F III,DATASET(ADD(SYS1.&SYSNAME..DATA))
MODIFIED  F RMF,F III,DATASET(ADD(SYS1.RMF3.DATA))
```

Defining parameters for Monitor I

RMF ships a default parmlib member (“ERBRMF00” on page 22) and an alternative one (“ERBRMF02” on page 23) to specify Monitor I gatherer options.

ERBRMF00

This is the default parmlib member for Monitor I gatherer sessions. It contains the options that RMF would default to anyway, if none were specified in a parmlib member. There are only two exceptions:

- The supplied parmlib member includes the option NOSTOP, whereas the RMF default is STOP(8H).
- The parmlib member includes NOOPTIONS instead of the RMF default OPTIONS. This suppresses the prompt for the operator to confirm the options, and so speeds up the start procedure.

The options are:

ERBRMF00

```
/*********************************************************************/
/* PART 1: MEASUREMENTS                                              */
/*********************************************************************/
  CACHE                   /* CACHE STATISTICS                       */
  CHAN                    /* CHANNEL STATISTICS                     */
  CPU                     /* CPU STATISTICS                         */
  CRYPTO                  /* CRYPTO HARDWARE                        */
  DEVICE(DASD)            /* DIRECT ACCESS DEVICES MEASURED         */
  DEVICE(NOTAPE)          /* NO TAPE DEVICES MEASURED               */
  DEVICE(NOCHRDR)         /* NO CHARACTER READER DEVICES MEASURED   */
  DEVICE(NOUNITR)         /* NO UNIT RECORD DEVICES MEASURED        */
  DEVICE(NOCOMM)          /* NO COMMUNICATION DEVICES MEASURED      */
  DEVICE(NOGRAPH)         /* NO GRAPHICS DEVICES MEASURED           */
  DEVICE(NONMBR)          /* NO SELECTION BY DEVICE NUMBERS         */
  DEVICE(NOSG)            /* NO SELECTION BY STORAGE GROUPS         */
NOENQ                     /* NO ENQUEUES MEASURED                   */
NOESS                     /* NO ENTERPRISE DISK SYSTEMS MEASURED    */
NOFCD                     /* NO FICON DIRECTOR MEASURED             */
  IOQ(DASD)               /* DASD I/O QUEUEING MEASURED             */
  IOQ(NOTAPE)             /* NO TAPE I/O QUEUEING MEASURED          */
  IOQ(NOCHRDR)            /* NO CHARACTER READER I/O QUEUEING       */
  IOQ(NOUNITR)            /* NO UNIT RECORD DEVICE I/O QUEUEING     */
  IOQ(NOCOMM)             /* NO COMMUNICATION I/O QUEUEING          */
  IOQ(NOGRAPH)            /* NO GRAPHICS DEVICE I/O QUEUEING        */
  IOQ(NONMBR)             /* NO SELECTIVITY BY LCU NUMBERS          */
  PAGESP                  /* PAGE DATASET STATISTICS                */
  PAGING                  /* PAGING DATA                            */
NOTRACE                   /* NO TRACE REPORT                        */
  VSTOR(S)                /* VIRTUAL STORAGE SUMMARY DATA           */
  WKLD                    /* WORKLOAD MANAGER DATA                  */
  NOVMGUEST               /* NO CPU DISPATCH TIMES FOR Z/VM GUEST   */
/*********************************************************************/
/* PART 2: TIMING                                                    */
/*********************************************************************/
```

```
  CYCLE(1000)              /* SAMPLE EVERY SECOND (1000 MSEC)      */
NOSTOP                     /* ACTIVE UNTIL OPERATOR ISSUES STOP    */
  SYNC(SMF)                /* USE INTVAL/SYNCVAL FROM SMFPRMXX     */
/*********************************************************************/
/* PART 3: REPORTING / RECORDING OF DATA                             */
/*********************************************************************/
NOOPTIONS                  /* OPTIONS NOT DISPLAYED, NO REPLY      */
  RECORD                   /* WRITE SMF RECORDS EVERY INTERVAL     */
NOREPORT                   /* NO WRITTEN REPORTS TO SYSOUT         */
  SYSOUT(A)                /* REPORTS TO CLASS A, IF REPORT        */
/*********************************************************************/
/* PART 4: USER EXITS                                                */
/*********************************************************************/
NOEXITS                    /* DO NOT TAKE USER EXITS               */
```

Note: If you miss gathering options for the coupling facility, for UNIX System Services or XCF, keep in mind that this data is gathered by Monitor III, and not by Monitor I.

ERBRMF02

This is the alternative parmlib member for Monitor I gatherer sessions. It contains options appropriate for monitoring of all resources in the system.

ERBRMF02

```
/*********************************************************************/
/* PART 1: MEASUREMENTS                                              */
/*********************************************************************/
  CACHE                    /* CACHE STATISTICS                     */
  CHAN                     /* CHANNEL STATISTICS                   */
  CPU                      /* CPU STATISTICS                       */
  CRYPTO                   /* CRYPTO HARDWARE                      */
  DEVICE(DASD)             /* DIRECT ACCESS DEVICES MEASURED       */
  DEVICE(TAPE)             /* TAPE DEVICES MEASURED                */
  DEVICE(CHRDR)            /* CHARACTER READER DEVICES MEASURED    */
  DEVICE(UNITR)            /* UNIT RECORD DEVICES MEASURED         */
  DEVICE(COMM)             /* COMMUNICATION DEVICES MEASURED       */
  DEVICE(GRAPH)            /* GRAPHICS DEVICES MEASURED            */
  DEVICE(NONMBR)           /* NO SELECTION BY DEVICE NUMBERS       */
  DEVICE(NOSG)             /* NO SELECTION BY STORAGE GROUPS       */
  ENQ(SUMMARY)             /* ENQUEUES MEASURED                    */
  ESS(LINK,RANK)           /* ENTERPRISE DISK SYSTEMS MEASURED     */
  FCD                      /* FICON DIRECTOR MEASURED              */
  IOQ(DASD)                /* DASD I/O QUEUEING MEASURED           */
  IOQ(TAPE)                /* TAPE I/O QUEUEING MEASURED           */
  IOQ(CHRDR)               /* CHARACTER READER I/O QUEUEING        */
  IOQ(UNITR)               /* UNIT RECORD DEVICE I/O QUEUEING      */
  IOQ(COMM)                /* COMMUNICATION I/O QUEUEING           */
  IOQ(GRAPH)               /* GRAPHICS DEVICE I/O QUEUEING         */
  IOQ(NONMBR)              /* NO SELECTIVITY BY LCU NUMBERS        */
  PAGESP                   /* PAGE DATASET STATISTICS              */
  PAGING                   /* PAGING DATA                          */
  TRACE(RCVUICA,END)       /* TRACE 'UIC AVERAGE'                  */
  TRACE(RCVCPUA,END)       /* TRACE 'CPU USAGE*16'                 */
  TRACE(RCVPTR,END)        /* TRACE 'PAGING RATE'                  */
  VSTOR(D)                 /* VIRTUAL STORAGE DETAIL DATA          */
  WKLD                     /* WORKLOAD MANAGER DATA                */
  NOVMGUEST                /* NO CPU DISPATCH TIMES FOR Z/VM GUEST */
/*********************************************************************/
/* PART 2: TIMING                                                    */
/*********************************************************************/
  CYCLE(250)               /* SAMPLE EVERY 250 MILLISECONDS        */
  STOP(8H)                 /* STOP AFTER 8 HOURS                   */
  SYNC(SMF)                /* USE INTVAL/SYNCVAL FROM SMFPRMXX     */
/*********************************************************************/
/* PART 3: REPORTING / RECORDING OF DATA                             */
/*********************************************************************/
  OPTIONS                  /* OPERATOR MAY EXAMINE/CHANGE OPTIONS  */
  RECORD                   /* WRITE SMF RECORDS EVERY INTERVAL     */
  REPORT(REALTIME)         /* WRITE REPORTS EACH INTERVAL          */
```

```
  SYSOUT(A)                /* REPORTS TO CLASS A, IF REPORT        */
/*****************************************************************/
/* PART 4: USER EXITS                                            */
/*****************************************************************/
NOEXITS                     /* DO NOT TAKE USER EXITS            */
```

Defining parameters for Monitor II

RMF ships a default parmlib member (“ERBRMF01” on page 24) and an alternative one (“ERBRMF03” on page 24) to specify Monitor II gatherer options.

ERBRMF01

This is the default parmlib member for Monitor II gatherer sessions. It contains the options that RMF would default to anyway, if none were specified in a parmlib member. There is only one exception; the supplied parmlib member includes the option STOP(30M), whereas the RMF default is STOP(10M). The options are:

ERBRMF01

```
/*****************************************************************/
/* PART 1: MEASUREMENTS                                          */
/*****************************************************************/
NOARD                       /* ADDRESS SPACE RESOURCE CONSUMPTION ? */
NOARDJ                      /* ARD REPORT FOR PARTICULAR JOB ?   */
  ASD                       /* ADDRESS SPACE STATE DATA ?        */
NOASDJ                      /* ASD REPORT FOR PARTICULAR JOB ?   */
NOASRM                      /* ADDRESS SPACE SRM DATA ?          */
NOASRMJ                     /* ASRM REPORT FOR PARTICULAR JOB ?  */
NOCHANNEL                   /* CHANNEL DATA ?                    */
NODEV                       /* DEVICE DATA ?                     */
NODEVV                      /* DEVICE DATA FOR SPECIFIC DEVICE ? */
NOIOQUEUE                   /* I/O QUEUING DATA ?                */
NOPGSP                      /* PAGE DATASET MEASUREMENTS ?       */
NOSENQ                      /* SYSTEM ENQUEUE CONTENTION ?       */
NOSENQR                     /* SYSTEM ENQUEUE RESERVE DATA ?     */
NOSPAG                      /* SYSTEM PAGING ACTIVITY ?          */
NOSRCS                      /* SYSTEM REAL STORAGE/CPU/SRM DATA ? */
/*****************************************************************/
/* PART 2: TIMING                                                */
/*****************************************************************/
  SINTV(30S)                /* SESSION INTERVAL = 30 SECONDS     */
  STOP(30M)                 /* STOP AFTER 30 MINUTES             */
/*****************************************************************/
/* PART 3: REPORTING / RECORDING                                 */
/*****************************************************************/
NODELTA                     /* TOTAL MODE                        */
NOOPTIONS                   /* NO OPERATOR DISPLAY, NO REPLY     */
  RECORD                    /* SMF RECORDING                     */
  REPORT(DEFER)             /* REPORTS PRODUCED AFTER SESSION END */
  SYSOUT(A)                 /* INTERVAL REPORTS TO CLASS A       */
/*****************************************************************/
/* PART 4: USER RECORDING/REPORTING                              */
/*****************************************************************/
NOUSER                      /* USER DATA ?                       */
```

ERBRMF03

This is the alternative parmlib member for Monitor II gatherer sessions. The contained options cause collection of data for all resources for a limited period.

The options are:

ERBRMF03

```
/*****************************************************************/
/* PART 1: MEASUREMENTS                                          */
```

```
/*********************************************************************/
  ARD                     /* ADDRESS SPACE RESOURCE CONSUMPTION ? */
NOARDJ                    /* ARD REPORT FOR PARTICULAR JOB ?      */
  ASD                     /* ADDRESS SPACE STATE DATA ?           */
NOASDJ                    /* ASD REPORT FOR PARTICULAR JOB ?      */
  ASRM                    /* ADDRESS SPACE SRM DATA ?             */
NOASRMJ                   /* ASRM REPORT FOR PARTICULAR JOB ?     */
  CHANNEL                 /* CHANNEL DATA ?                       */
  DEV                     /* DEVICE DATA ?                        */
NODEVV                    /* DEVICE DATA FOR SPECIFIC DEVICE ?    */
  IOQUEUE                 /* I/O QUEUING DATA ?                   */
  PGSP                    /* PAGE DATASET MEASUREMENTS ?          */
  SENQ                    /* SYSTEM ENQUEUE CONTENTION ?          */
  SENQR                   /* SYSTEM ENQUEUE RESERVE DATA ?        */
  SPAG                    /* SYSTEM PAGING ACTIVITY ?             */
  SRCS                    /* SYSTEM REAL STORAGE/CPU/SRM DATA ?   */
/*********************************************************************/
/* PART 2: TIMING                                                   */
/*********************************************************************/
  SINTV (30S)             /* SESSION INTERVAL = 30 SECONDS        */
  STOP (1H)               /* STOP AFTER 1 HOUR                    */
/*********************************************************************/
/* PART 3: REPORTING / RECORDING                                    */
/*********************************************************************/
  DELTA                   /* PRESENT DATA AS INTERVAL DELTAS      */
  OPTIONS                 /* OPERATOR MAY EXAMINE/CHANGE OPTIONS  */
  RECORD                  /* SMF RECORDING                        */
  REPORT (DEFER)          /* REPORTS PRODUCED AFTER SESSION END   */
  SYSOUT(A)               /* INTERVAL REPORTS TO CLASS A          */
/*********************************************************************/
/* PART 4: USER RECORDING/REPORTING                                 */
/*********************************************************************/
NOUSER                    /* DO NOT COLLECT USER DATA             */
```

Defining parameters for Monitor III

RMF ships a default parmlib member (“ERBRMF04” on page 25) to specify Monitor III gatherer options.

ERBRMF04

This is the default parmlib member for Monitor III data gatherer sessions. There is no IBM supplied alternative member for this gatherer.

The options specified in ERBRMF04 are:

ERBRMF04

```
  CYCLE(1000)             /* SAMPLE EVERY SECOND (1000 MSEC)      */
  DATASET(STOP)           /* NO DATASET SUPPORT                   */
  DATASET(NOSWITCH)       /* APPEND TO LAST NON-FULL DATASET      */
  DATASET(WHOLD(128))     /* CONTROLS BUFFER PAGES IN STORAGE     */
  MINTIME(100)            /* LENGTH OF MINTIME                    */
  NOOPTIONS               /* DO NOT DISPLAY OPTIONS               */
  RESOURCE(*JES2,JES2)    /* SPECIFIES JES STARTED TASK NAME      */
  NOSTOP                  /* RUN UNTIL OPERATOR ISSUES STOP       */
  SYNC(00)                /* MINTIME SYNCHRONIZATION              */
  SYSOUT(A)               /* MESSAGES TO SYSOUT CLASS A           */
  WSTOR(128)              /* SIZE OF INSTORAGE BUFFER (IN MB)     */
  MASTER                  /* SYSTEM IS ELIGIBLE FOR MASTER        */
  ZIIPUSE                 /* PARTIAL USE OF zIIP ENGINES          */
  IOSUB                   /* I/O SUBSYSTEM GATHERING ACTIVE       */
  CFDETAIL                /* COUPLING FACILITY DETAILS            */
  CACHE                   /* ACTIVATE CACHE GATHERING             */
  VSAMRLS                 /* ACTIVATE VSAM RLS GATHERING          */
  OPD                     /* ACTIVATE OMVS PROCESS DATA GATHERING */
  PCIE                    /* ACTIVATE PCIE DATA GATHERING         */
  EADM                    /* ACTIVATE EADM (AKA SCM) DATA GATH.   */
  ZFS                     /* ZFS DATA GATHERING                   */
```

```
  NOSGSPACE                  /* NO STORAGE GROUP SPACE GATHERING     */
  NOLOCK                     /* NO LOCK DATA GATHERING               */
```

Part 3. Operation

Operation is what you have to do at the system console to start the RMF control session and certain monitor sessions. The default data sets and monitor options should already have been defined by the administrator.

An operator can override the default monitor options. It is best to do this in agreement with the performance analyst who will be evaluating the data that RMF gathers.

What Operation Involves

Using system commands, the operator can:

- Start and stop the RMF control session
- Start and stop individual background sessions
- Specify monitor options that are to be valid for the session, as opposed to your system's default options; or change options during a monitor session.
- Influence the SMF data buffer, in which RMF data is stored

The monitor options which you can specify on the system commands are described in detail in Part 4, “Data gathering reference,” on page 47 and Reporting reference in *z/OS RMF Reporter User's Guide*.

Chapter 3. Starting and stopping RMF

This information unit explains:

- How to start the RMF control session:

  Use the system command START to start the RMF control session, or to start both the control session and a Monitor I session. After you have started the control session, you can start all monitor sessions from the console, except Monitor II and Monitor III TSO/E sessions.
- How to specify the SMF buffer:

  RMF data gatherers write data as SMF records that can be stored in an in-storage, wrap-around SMF buffer for further processing.
- How to stop RMF.

Starting RMF

Enabling RMF:

RMF is an optional element of z/OS. It is present in the system, whether you have purchased it or not. If you have not specifically ordered RMF, it is disabled. This start procedure will not work, and you will receive the message:

ERB111I RMF IS NOT ENABLED TO RUN ON THIS SYSTEM

It is the task of the system administrator to see to it that RMF, if licensed, is enabled to run.

The system command START invokes the RMF cataloged procedure, and you can override specifications in the procedure JCL statements with specifications on the START command.

The syntax of the START command for RMF is:

```
{START} RMF,,,[parm]
{S    }
```

parm
Can serve the following purposes:

- specifying the SMF buffer options (see "Controlling the SMF buffer" on page 31)
- specifying the Monitor I gatherer session options (see Chapter 6, "Long-term data gathering with Monitor I," on page 49)

  If you specify options, each must have the form:

  ```
  option [(value)] for example: DEVICE(DASD)
  ```
- providing automatic sysplex-wide management for the Distributed Data Server (see Starting the Distributed Data Server in *z/OS RMF Reporter User's Guide*)

Multiple options must be enclosed in parentheses and separated by commas, for example:

```
(DEVICE(DASD),CYCLE(500),DDS)
```

By default, Monitor I is started along with RMF. If options are specified for `parm`, they will be used. To start a Monitor I session using options from the default Parmlib member ERBRMF00 or program defaults, omit this parameter. See Chapter 5, "How RMF processes session options," on page 41.

If you do not want to start the Monitor I session, specify only **NOZZ** for `parm`.

Note: RMF can not run in reusable address spaces. It is not possible to specify REUSASID=YES on the START command.

Examples:

- To start the RMF control session only, issue the system command:

```
START RMF,,,NOZZ
```

- To start both RMF control and Monitor I sessions, specify:

```
START RMF
```

- To start RMF with a Monitor I session and the Distributed Data Server, issue the command:

```
START RMF,,,DDS
```

- To start both RMF control session and Monitor I, and specify options, issue the command:

```
START RMF,,,(DEVICE(DASD),CYCLE(500))
```

- To start RMF with a Monitor I session and an SMF buffer of 32 megabytes in which SMF record types 72 to 74 are to be stored, specify:

```
START RMF,,,(SMFBUF(RECTYPE(72:74)))
```

Starting RMF in the sysplex

You have to start the RMF control session and the data gatherer sessions separately on each system of the sysplex, if you want sysplex-wide reports. The reporting, however, can be done on one system.

We strongly recommend that you start RMF on all systems in a sysplex with the same options. This is essential for later sysplex reporting. The easiest way to do this is by using the ROUTE command, as shown in this example:

```
RO   *ALL,S RMF
```

Starting RMF without JES

You can run the RMF data gatherers on a system on which a job-entry subsystem (JES2 or JES3) is not active, if you take the following steps:

1. Preallocate the RMF message data sets.

   In the RMF procedure in SYS1.PROCLIB, include DD statements for the message data sets. If you wish, you can specify "DD DUMMY". Allocate the following DDNAMES:

   MFMESSGE
   for RMF general messages

   RMFSCZZ
   for Monitor I session messages

   RMFSCIII
   for Monitor III session messages

   RMFSCxx
   for Monitor II background messages from session xx

   If you intend to start the Monitor III data gatherer, also preallocate the DDNAME RMFM3III in the RMFGAT procedure in SYS1.PROCLIB (see “Setting up the Monitor III gatherer session RMFGAT” on page 16).

   Since RMF is running with NODSI setting in its PPT entry, be aware that RMF does not hold any data set ENQUEUE (major name=SYSDSN, minor name=dsname) for the data sets allocated by the RMF procedure. A missing ENQUEUE can mislead functions like HSM, that rely on ENQUEUE for data set protection.
2. Specify SUB=MSTR on the START command.

Enter the START command in the following format:

```
{START} RMF,,,[parm],SUB=MSTR
{S    }
```

parm
Other options as described in “Starting RMF” on page 29.

SUB=MSTR
Use this specification if JES is not active on your system, and you want to run the RMF data gatherers.

3. Suppress the printing of reports.

   Start the gatherer sessions in the normal way (see “Starting RMF sessions” on page 35) but be sure to specify the NOREPORT option for both Monitor I and Monitor II background sessions.

Stopping RMF

The system command STOP ends the RMF control session and all active gatherer and background sessions. Any active Monitor II and Monitor III TSO/E sessions remain active. RMF issues a message informing you that RMF has stopped. For information about stopping individual sessions, see “Stopping RMF sessions” on page 39. The syntax of the STOP command for RMF is:

```
{STOP}    RMF
{P   }
```

Controlling the SMF buffer

RMF data gatherers write data to SMF records, from which the Postprocessor can extract the information you request. The SMF records can be written to SMF data sets or to SMF log streams, but they can also be written to an in-storage, wrap-around SMF buffer (see “Accessing data across the sysplex” on page 6). You can control the size of this buffer and the SMF record types that RMF writes to it, using the SMFBUF option.

The RMF default values for the SMF wrap-around buffer are:

- a size of 32 megabytes
- collection of SMF record types 70 to 78, including all subtypes

You can override these values by specifying the SMFBUF option in any of three ways. In each case, the keywords SPACE and RECTYPE with the desired values can be specified:

1. By specifying SMFBUF in the PARM field of the cataloged procedure which starts the RMF control session (see “Setting up the RMF control session including Monitor I and Monitor II” on page 15). This overrides the RMF default values.
2. By specifying the SMFBUF option on the system command START for the RMF control session. This overrides any PARM specification, and the RMF defaults.
3. By specifying the SMFBUF option on the system command MODIFY for the RMF control session. This overrides any specifications on the START command, or in the cataloged procedure, and the RMF defaults.

The format of the SMFBUF option is:

```
NOSMFBUF
or
SMFBUF[(
[SPACE(size{K|M|G}
[,FORCE])]
[,RECTYPE(rtype)])]
```

The default is **NOSMFBUF**.

size
Is a positive integer specifying the size of the buffer, and K, M and G stand for kilobytes, megabytes and gigabytes, respectively.

The minimum size of the data buffer is 1M or 1024K, the maximum size is 2G. If SMFBUF is specified without `size`, the size of the buffer defaults to 32M.

FORCE
As a keyword on the SPACE parameter is meaningful only on the MODIFY command, not on START or in the cataloged procedure. It causes the size of an existing SMF data buffer to be adjusted immediately. If FORCE is not specified, the data buffer size is adjusted during the next wrap-around interval, which depends on the current size of the data buffer.

When you reduce the size of an already active SMF buffer, bear in mind that FORCE will cause a loss of any data stored at the upper end of the old buffer.

rtype
Specifies the SMF record type or types to be stored in the buffer. Valid values are:

- A decimal number from 0 to 255, inclusive, denoting an SMF record type. You can follow each record type with a list of subtypes in parentheses.
- Two such numbers, separated by a colon(:), denoting a range of SMF record types. No subtypes can be specified, in this case.

If you specify a record type without a subtype list, or a record type range, all subtypes of the specified record type or types are stored in the data buffer.

Note: SMF records type 79 subtype 15 (for Monitor II IRLM long lock reporting) will be written only if you define this explicitly, for example

`SMFBUF(RECTYPE(0:78,79(1:15)))`

If you omit `rtype`, the default value used is `70:78`.

SMFPRMxx in SYS1.PARMLIB
To write SMF records type 79 subtype 15 (for Monitor II IRLM long lock reporting), exits `IEFU83, IEFU84, IEFU85` need to be defined, for example:

```
SYS(.....,EXITS(IEFU83,IEFU84,IEFU85,....
SUBSYS(STC,EXITS(IEFU83,IEFU84,IEFU85,............
```

The defaults mean that SMFBUF without options in the cataloged procedure or on the START command is equivalent to:

```
SMFBUF(SPACE(32M),RECTYPE(70:78))
```

If you specify SMFBUF without options on the MODIFY command, RMF displays the current options, or tells you if the data buffer is not active.

The values specified on a system command override any SMFBUF option in the RMF cataloged procedure.

Examples: Assume you have included in your RMF cataloged procedure:

```
//EFPROC EXEC PGM=ERBMFMFC,REGION=256M,
//       PARM='SMFBUF(SPACE(40M),RECTYPE(70:79))'
```

This will be your system's standard SMF buffer definition. SMF records of types 70 to 79 inclusive will be stored in a 40-megabyte wrap-around buffer.

To alter the record types for one RMF control session, use the START command, for example:

```
S RMF,,,(SMFBUF(RECTYPE(72(1,2,3))))
```

This leaves the size of the wrap-around buffer unchanged, but causes only SMF records of type 72, subtypes 1, 2 and 3 to be stored in it.

During the RMF control session, you can alter the size of the SMF wrap-around buffer without affecting the record types to be collected. Use the MODIFY command to reduce the size of the buffer, for example:

```
F RMF,SMFBUF(SPACE(16M))
```

To make SMF records type 104 available in the SMF buffer, you must specify SMF record type 104, optionally reduced to required subtypes, with the SMFBUF parameter, for example:

```
SMFBUF(RECTYPE(70:78,79(1:15),104(1:12)))  /* for AIX on System p */
SMFBUF(RECTYPE(70:78,104(20:31)))          /* for Linux on System x */
SMFBUF(RECTYPE(70:79,104(40:53)))          /* for Linux on System z */
```

Chapter 4. Starting and controlling Monitor sessions

This information unit explains how to start and stop RMF sessions, specify and modify session options, and display status for the following:

- Monitor I session
- Monitor II background sessions
- Monitor III gatherer session
- RMF Client/Server sessions (RMFCS)

You can find details of the options and commands for all RMF sessions in Part 4, “Data gathering reference,” on page 47 and in Reporting reference in *z/OS RMF Reporter User's Guide*.

Specifying session options

When starting or modifying the sessions described in this topic, you can specify options on the system command MODIFY. However, you need not do this for every session, if you have specified your own installation default options elsewhere. You can do this in:

- The PARM field of the EXEC statement in the RMF cataloged procedure (Monitor I session only. See “Setting up the RMF control session including Monitor I and Monitor II” on page 15 for more details).
- The RMF Parmlib member, or other equivalent data set member containing session options. See “Storing gatherer options” on page 19 for more details.

If you do not specify an option in either the MODIFY command, the PARM field or the Parmlib member, RMF uses a program default. From the various specifications, RMF forms a list of options for the session. How it does this is described in Chapter 5, “How RMF processes session options,” on page 41.

Conflicts between session options

Some options cannot be used concurrently, and may cause conflicts. 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 for each monitor are discussed in Part 4, “Data gathering reference,” on page 47 and in Reporting reference in *z/OS RMF Reporter User's Guide*.

Starting RMF sessions

Session commands are issued as parameters on the **system** command **MODIFY**. Only one Monitor I session can be active at any particular time; up to 32 non-interactive Monitor II sessions can be active concurrently.

RMF provides a cataloged procedure which starts a Monitor III data gatherer session, as described in “Setting up the Monitor III gatherer session RMFGAT” on page 16. It is invoked in response to the session command START. If you want to modify the JCL by adding parameters, you must do so before starting the session. See “Starting a Monitor III gatherer session” on page 36.

Starting a specific Monitor

Once you have started the RMF control session, use the system command MODIFY to pass the session command START to it. The syntax of the START session command is:

```
{MODIFY} RMF,{START} session-id [,parm]
{F     }     {S    }
```

session-id
Identifies which monitor session to start:

- **ZZ** for Monitor I
- **Two alphameric characters** for a Monitor II background session, but not ZZ.
- **III** for the Monitor III gatherer session

For the Monitor II sessions, of which you can start several at a time, `session-id` distinguishes the various sessions. Use this session-ID on all session commands for that particular session. The session-ID also appears in all RMF messages about that session.

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

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

If you specify multiple options, they must be separated by commas.

For guidelines on specifying options, see the sections on starting the respective sessions below.

If you do not specify session options here, RMF takes all options from the Parmlib member and program defaults. See Chapter 5, “How RMF processes session options,” on page 41 for information about how RMF processes options when you start an RMF session.

Starting a Monitor I session

The value of `session-id` for Monitor I is always **ZZ**. If you start the Monitor I session when you start RMF, ZZ is automatically assigned as the session-ID.

If JES is not active in your system, and you have started RMF with the SUB=MSTR option, you must specify the NOREPORT option when starting this gatherer. This and other options and values that you can specify for `parm` are listed in Chapter 6, “Long-term data gathering with Monitor I,” on page 49.

Example: To start the Monitor I session, specifying that processor activity is not to be measured, and take all other options from other sources, issue the command:

```
MODIFY RMF,START ZZ,NOCPU
```

Starting a Monitor II background session

The value of `session-id` for a Monitor II background session can be any two-character alphameric value except ZZ.

If JES is not active in your system, and you have started RMF with the SUB=MSTR option, you must specify the NOREPORT option when starting this gatherer. This and other options and values that you can specify for `parm` are listed in Details of report commands in *z/OS RMF Reporter User's Guide*.

Example:

- To start a Monitor II background session when all options are to be taken from the program defaults, issue the command:

  ```
  MODIFY RMF,START AB
  ```

- 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 Parmlib member ERBRMF07, issue the command:

  ```
  MODIFY RMF,START BB,MEMBER(07),REPORT(REALTIME)
  ```

Starting a Monitor III gatherer session

The value for `session-id` is always **III**. Specify this in the START command.

RMF invokes the Monitor III cataloged procedure (RMFGAT) in response to the Monitor III gatherer session command START (see "Setting up the Monitor III gatherer session RMFGAT" on page 16). If you want to modify the JCL procedure by specifying parameters, you must do so before starting the session.

The options and values that you can specify for `parm` are listed in Chapter 8, "Short-term data gathering with Monitor III," on page 73.

Examples:

- To start a data gatherer session with all options taken from the Parmlib member and the program defaults, issue the following command:

```
MODIFY  RMF,START III
```

- To start a data gatherer that is to sample data at a 2000 millisecond cycle, combine samples after a 300 second interval, and run for 12 hours, issue the following command:

```
MODIFY RMF,START III,CYCLE(2000),MINTIME(300),STOP(12H)
```

In case of a time change, for example, for a switch to daylight saving time, RMFGAT will be stopped and restarted automatically to reflect the correct time in the sampled records.

Starting an RMF client/server session (RMFCS)

RMF Client/Server Enabling (RMFCS) is a concept that supports performance management for z/OS systems without an active TSO/TCAS subsystem on the host. RMFCS allows you to establish as many sessions as you want with any systems in your network that have an APPC or TCP/IP connection configured to your workstation. This is possible with all operating systems that support the ISPF Client/Server component.

Within one session, you can have up to 32 active windows by using the ISPF/SPLIT function, which allows 32 logical screens. Each SPLIT creates a new window, and you can toggle through your windows by using the SWAP function, which shifts the focus to the next window.

This way, RMFCS combines the advantages of a single point of control for z/OS performance management with a state-of-the-art user front end.

You can access RMF Monitor II and Monitor III reports with RMFCS by exploiting the ISPF Batch GUI feature.

Either start procedure RMFCSC by commands shown below, or add the commands to the appropriate Parmlib member during IPL of the system:

Example:

To start an RMFCS for TSO-users USER#1 and USER#2, issue the command:

```
S RMFCSC,HLQ=USER#1
S RMFCSC,HLQ=USER#2
```

You can find details about RMFCS in RMF Client/Server (RMFCS) enabling in *z/OS RMF Reporter User's Guide*.

Modifying RMF session options

You can modify the options in effect for the sessions described in this chapter, using the MODIFY command. A changed option remains in effect until the session ends or you issue the MODIFY command to change the option again. The syntax of the MODIFY command is:

```
{MODIFY} RMF,{MODIFY} session-id[,parm]
{F     }     {F     }
```

session-id
The identifier you specified on the session command START.

parm
The options for the rest of the session. Each option has the form:

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

If you specify multiple options, you must separate them by commas.

The options that you can specify are the same as on the session command START.

For information about how RMF processes options when you modify session options, see Chapter 5, "How RMF processes session options," on page 41.

Examples:

- Monitor I session:

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

  ```
  MODIFY RMF,MODIFY ZZ,CPU
  ```

- Monitor II background session with the session-ID AB:

  To modify the options to add printed output to SMF record output (NOREPORT and RECORD in effect), enter the command:

  ```
  MODIFY RMF,MODIFY AB,REPORT(DEFER)
  ```

- Monitor III gatherer session:

  To modify the NOSTOP option to STOP (after a duration of four hours) and change the time interval to 200 seconds, issue the command:

  ```
  MODIFY RMF,MODIFY III,STOP(4H),MINTIME(200)
  ```

Note: Modifying Monitor I options using the MODIFY command will cause a reinitialization of the complete Monitor I ZZ session.

Displaying RMF status

To determine what sessions are active and what options are in effect, you can display the RMF status from the operator console:

```
                          {ACTIVE    }
{MODIFY} RMF, {DISPLAY} {session-id}
{F     }      {D      } {ALL       }
```

ACTIVE
Specifies that the session-IDs of all active sessions are to be displayed. This is the default value.

session-id
Specifies the session-ID of a particular session. The options for the named session are displayed.

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

Examples:

- To display the session identifiers of all active background sessions, issue the command:

  ```
  MODIFY RMF,DISPLAY ACTIVE
  ```

  or

```
F RMF,D
```

to use the shortest form.

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

```
F RMF,D ZZ
```

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

```
F RMF,D ALL
```

- To display the console output produced for a particular Monitor III data gatherer session, issue the command:

```
F RMF,D III
```

- To display the SMFBUF option, issue the command:

```
F RMF,SMFBUF
```

Stopping RMF sessions

You can end sessions in three ways:

- By issuing the system command STOP, which stops all active background sessions. See “Stopping RMF” on page 31.
- By specifying a time value in the STOP option for a specific session. See Part 4, “Data gathering reference,” on page 47.
- By issuing a STOP session command to stop a specific session. All other active sessions continue processing. See “Stopping a specific session” on page 39.

Stopping a specific session

You can end any active session with the command:

```
{MODIFY} RMF, {STOP} session-id
{F     }      {P   }
```

session-id
The identifier assigned on the START command for your session.

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

Note that stopping Monitor I influences other monitors that are using data gathered by Monitor I.

Examples:

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

```
MODIFY RMF,STOP ZZ
```

- To stop a Monitor II background session with an identifier AB, issue the command:

```
MODIFY RMF,STOP AB
```

- To stop the Monitor III gatherer, while allowing all other active sessions to continue processing, issue the command:

```
MODIFY RMF,STOP III
```

Chapter 5. How RMF processes session options

RMF processes session options from various sources in a certain order to create a list of options for a non-interactive session. RMF uses a list of options to control each non-interactive session:

- Monitor I session
- Monitor II background session
- Monitor III gatherer session

RMF processes session options whenever you use:

- A START command to start Monitor I when you start RMF
- A session command START to start non-interactive sessions
- A session command MODIFY to modify non-interactive session options

This information unit describes how RMF processes session options in all of these situations.

When you start an RMF session

When you start a non-interactive session from the operator console, RMF processes the options from the following sources, listed here in order of priority:

1. The parm field of the START session command (highest priority).

   The options you specify here override any others.
2. The PARM field in the EXEC statement of the RMF cataloged procedure.

   This source is relevant only when you use the system command START to start Monitor I along with the RMF control session.
3. The specified Parmlib members.

   If you include the option MEMBER in the START command or in the RMF cataloged procedure, the options in the specified Parmlib member are taken next.

   If you specify more than one member, RMF gives precedence to the options in the member specified first in the list. For example, if you specify MEMBER(02,07), RMF first notes the options from ERBRMF02, then processes those from ERBRMF07. In case of conflicts, RMF uses the options from ERBRMF02. This means that, if ERBRMF02 specifies ENQ(DETAIL) and ERBRMF07 specifies ENQ(SUMMARY), RMF establishes ENQ(DETAIL) for the session.

   The default Parmlib member is not used if the MEMBER option is in effect.
4. The RMF default Parmlib members.

   If you do not specify the MEMBER option in any of the above places, RMF uses the default Parmlib members. RMF establishes options from the default Parmlib members only if they were not specified in any of the higher-priority places listed above.
5. Program defaults (lowest priority).

   RMF fills in those options not specified anywhere else with a program default. The program defaults for non-interactive session options are described in each respective chapter.

If RMF encounters any conflicting options while processing the session options, it chooses the value specified in the higher-priority source, and issues a warning message. For example, RMF might detect the Monitor II background session options RECORD on the START command and NORECORD in a Parmlib member. Since RMF detected RECORD higher in the priority list, it takes that value.

If RMF detects invalid option values, it ignores them and uses the next valid value specified in priority source.

If RMF does not find any errors, it issues a informational message indicating that the session is active, and begins session processing.

Displaying a list of options in effect for a session

If RMF detects any errors while processing session options, it displays a list of options in effect for a non-interactive session to the operator console, and issues a message. You can respond to the message by correcting the invalid options or specifying additional options. You can display a list of options in effect for a non-interactive session at any time by:

- Issuing the DISPLAY session command from the operator console. For information about issuing a DISPLAY session command, see "Displaying RMF status" on page 38.
- Specifying the session option OPTIONS.

Examples

This section shows how RMF processes session options when you start non-interactive sessions.

When you start a Monitor I session

Assume that you start a Monitor I session along with the RMF control session, using the following system command:

```
START RMF,,,(MEMBER(10),CYCLE(1000),DEVICE(COMM))
```

From the options specified in the START system command, RMF creates the following option list for the session:

```
CYCLE(1000)
DEVICE(COMM)
```

RMF processes the MEMBER(10) option after it processes all other options specified in the START system command. Member ERBRMF10 contains the following options:

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

After processing ERBRMF10, the option list for the session is now:

```
CYCLE(1000)
DEVICE(COMM,NOUNITR,TAPE)
NOEXITS
```

RMF processes the next option source, the PARM= field of the RMF cataloged procedure. The START system command invokes the following user-modified cataloged procedure:

```
//IEFPROC EXEC PGM=ERBMFMFC,REGION=256M,
//             PARM='CYCLE(2000),DEVICE(NOTAPE,DASD),
//             MEMBER(02)'
```

RMF processes the options specified in the PARM= field of the RMF cataloged procedure and the option list is now:

```
CYCLE(1000)
DEVICE(COMM,NOUNITR,TAPE,DASD)
NOEXITS
```

RMF ignores CYCLE(2000) and DEVICE(NOTAPE) because these options have been filled in by a higher-priority source.

RMF processes the MEMBER(02) option after it processes all other options specified in the START system command. Member ERBRMF02 contains the following options:

```
OPTIONS
NOPAGESP
EXITS
```

RMF processes the member, and the option list is now:

```
CYCLE(1000)
DEVICE(COMM,NOUNITR,TAPE,DASD)
NOEXITS
OPTIONS
NOPAGESP
```

RMF ignores EXITS specified in member ERBRMF02 because it already filled those in from a higher priority source. RMF adds NOPAGESP from ERBRMF02.

Because not all options have been filled in, RMF uses program defaults to complete the option list.

When you start a Monitor II background session

Assume that the operator issued the following START command to start a Monitor II background session:

```
MODIFY RMF,START AB,DELTA,SINTV(30),MEMBER(07)
```

RMF 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. RMF would not use SINTV(10) 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, RMF proceeds to the IBM supplied program defaults. (These defaults are indicated in the discussion of each option in Details of report commands in *z/OS RMF Reporter User's Guide*.)

After adding the defaults, RMF builds a complete list of session options:

```
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 you modify session options

When you use the MODIFY session command to modify the options for a non-interactive session, RMF processes the options in a different priority order than when you start a non-interactive session. RMF starts with the list of options previously established and uses the input sources to **override** any previously established option.

The input sources have the following order of priority:

1. **The options field of the session command MODIFY**.

   Any options you specify here override and replace any options in effect prior to the MODIFY command.

2. **RMF Parmlib members, in a left to right order**

   If you include a MEMBER option in the options field of the MODIFY command, any options specified in the member override any options specified previously in the MODIFY command.

   When you specify more than one member, RMF processes the members in left to right order; the rightmost member overriding any corresponding options from a previously-processed member.

   Example:

   If you specify MEMBER(03,07) on a MODIFY command, RMF generates the member names ERBRMF03 and ERBRMF07 and proceeds as follows:

   - Take the options from ERBRMF03 first. ERBRMF03 specifies NOASD, so the merge process places NOASD in the list of session options.
   - Now take the options from member ERBRMF07. ERBRMF07 specifies ASD, so the merge process places ASD in the list of session options.

   ASD overrides the previously-established NOASD, and ASD is valid for the session.

RMF responds to errors in a MODIFY session command in the same way as in a START session command.

Examples

This section shows how RMF processes session options for non-interactive sessions when you use a MODIFY session command.

When you modify Monitor I 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,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,MODIFY ZZ,STOP(40M),MEMBER(10)
```

RMF:

1. Merges the input option from the command and replaces NOSTOP in the current option list with STOP(40M).

2. Merges 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.

Monitor II background session

Assume that the options for a currently-active Monitor II background 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,MODIFY AB,ASD,NOSPAG,STOP(40M)
```

RMF modifies the options as you want.

If, however, member ERBRMF09 includes the options:

```
ASD
NOSPAG
NOSTOP
```

and you issue the command:

```
MODIFY RMF,MODIFY AB,STOP(40M),MEMBER(09)
```

RMF:

1. Replaces NOSTOP in the current option list with STOP(40M).
2. Reads ERBRMF09, compares options from that member with the current options list, and replaces NOASD with ASD, SPAG with NOSPAG, and STOP(40M) with NOSTOP.

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

When you modify Monitor III data gatherer options

Assume that the options for a currently active session include CYCLE(500), MINTIME(50) and NOSTOP, and that you want to modify these options to CYCLE(1000), MINTIME(200) and STOP(40M).

If you issue the command:

```
MODIFY  RMF,MODIFY III,CYCLE(1000),MINTIME(200),STOP(40M)
```

the options will be modified as you want.

If, however, member ERBRMF10 includes the options:

```
CYCLE(1000)
MINTIME(200)
NOSTOP
```

and you issue the command:

```
MODIFY  RMF,MODIFY III,STOP(40M),MEMBER(10)
```

RMF:

1. Merges the input option from the command and replaces NOSTOP in the current option list with STOP(40M).
2. Merges the options from ERBRMF10 with the current options list, replacing CYCLE(500) with CYCLE(1000), MINTIME(50) with MINTIME(200) and STOP(40M) with NOSTOP.

In this particular case, the desired STOP(40M) option is not currently in effect. This particular command did not achieve the expected results because any option in a member will override both the corresponding current option and the corresponding option specified on the MODIFY session command.

To modify the NOSTOP option of an active data gatherer session to STOP (after a duration of four hours) and change the time interval to 200 seconds, issue the command:

```
MODIFY RMF,MODIFY III,STOP(4H),MINTIME(200)
```

Part 4. Data gathering reference

This part deals with the RMF data gathering capabilities, and with how to control them:

- long-term gathering with Monitor I
- snapshot gathering with Monitor II
- short-term gathering with Monitor III

All the options and commands you need are described fully in the appropriate chapters.

Chapter 6. Long-term data gathering with Monitor I

This information unit describes the Monitor I gatherer session options in alphabetical order. The program defaults are underscored where appropriate.

You can specify Monitor I session options in:

- the **parm** field of the START command that starts the session (see "Starting a specific Monitor" on page 35)
- the PARM field of the EXEC statement in the RMF cataloged procedure (see "Setting up the RMF control session including Monitor I and Monitor II" on page 15)
- the RMF Monitor I parmlib member ERBRMF00 (see "ERBRMF00" on page 22)

RMF merges the input to a final set of options for the session. See Chapter 5, "How RMF processes session options," on page 41 for details.

Summary of session options

Monitor I creates SMF records type 70 — 78. You can find an overview in Table 2 on page 5.

Table 3 on page 49 gives a summary of the available options, grouped by purpose. There are options for specifying:

- What activities to monitor
- The time-frame for monitoring them
- What reports to produce
- Environmental information

Table 3. Summary of Monitor I Session Options

Option	Description	Details on
Activity Options		
CACHE/NOCACHE	Cache activity	"CACHE" on page 51
CHAN/NOCHAN	Channel path activity	"CHAN" on page 52
CPU/NOCPU	Processor activity	"CPU" on page 52
CRYPTO/NOCRYPTO	Cryptographic hardware activity	"CRYPTO" on page 52
DEVICE(type)/NODEVICE	Device activity	"DEVICE" on page 52
{(SUMMARY)} ENQ{(DETAIL[,majorname[,minorname]])}/ NOENQ	Enqueue contention activity	"ENQ" on page 54
ESS[(opt_list)]/NOESS	Enterprise Disk System statistics	"ESS" on page 55
FCD/NOFCD	FICON® director activity	"FCD" on page 56

Table 3. Summary of Monitor I Session Options (continued)

Option	Description	Details on
IOQ (opt_list)/NOIOQ	I/O queuing activity	"IOQ" on page 57
PAGESP/NOPAGESP	Page data set activity	"PAGESP" on page 60
PAGING/NOPAGING	System paging activity	"PAGING" on page 60
TRACE(variable[,opt_list])/NOTRACE	Trace variables for the Trace Activity report	"TRACE" on page 62
{(S)} VSTOR{(D [,jobname1,jobname2,...])}/ NOVSTOR	Virtual storage activity	"VSTOR" on page 67
WKLD/NOWKLD	Workload activity	"WKLD" on page 68
VMGUEST/NOVMGUEST	Processor activity	"VMGUEST" on page 68
Time-frame Options		
{(1000)} CYCLE{(nnn)}	The length of the cycle at the end of which RMF makes sampling observations	"CYCLE" on page 52
({30M}) INTERVAL({nnn[M]})	The length of the reporting interval in minutes used in combination with the options SYNC(RMF,mm), SYNC(RMF,mmM), or NOSYNC	"INTERVAL" on page 56
[M] STOP (value[H])/NOSTOP	Desired duration of the Monitor I session, in minutes (M), or hours (H)	"STOP" on page 61
{(SMF)} {(RMF,mm)} SYNC {(RMF,mmM)}/NOSYNC	Interval synchronization with the SMF or the RMF interval synchronization on the minute	"SYNC" on page 61
Reporting Options		
{OPTIONS}/{NOOPTIONS} {OPTN } {NOOPTN }	Option list for the session to be printed at the operator console	"OPTIONS" on page 59
RECORD/NORECORD	Specifies whether measured data is to be written to SMF records	"RECORD" on page 60
{(REALTIME)} REPORT{(DEFER)}/NOREPORT	Specifies production of printed interval reports of measured data	"REPORT" on page 60
SYSOUT(class)	SYSOUT class to which the formatted printed reports are directed	"SYSOUT" on page 62
Environment Options		
EXITS/NOEXITS	User exit routines to be executed during session processing to gather or report additional data	"EXITS" on page 55
MEMBER(list)	Parmlib member containing Monitor I session options	"MEMBER" on page 58

Default session options

Here are the options that take effect by default. You need only specify contradictory ones:

Table 4. Monitor I Default Session Options

Default Option	Description
CACHE	Measures cache activity
CHAN	Measures channel path activity
CPU	Measures processor activity
CRYPTO	Measures cryptographic hardware activity
CYCLE(1000)	Takes sample measurements once a second (1000 milliseconds)
DEVICE(DASD)	Measures DASD activity (not other classes of device)
IOQ(DASD)	Measures I/O queuing activity on logical control units for DASD
NOESS	Does not measure Enterprise Disk System statistics
NOFCD	Does not measure FICON director activity
PAGESP	Measures page data set activity
PAGING	Measures system paging activity
RECORD	Writes the measured data to SMF records
VSTOR(S)	Measures virtual storage activity for summary reports
WKLD	Measures system workload
INTERVAL(30M)	Combines data every 30 minutes (value is ignored for SYNC(SMF))
NOENQ	Does not measure contention activity
NOEXITS	Executes no user exits when gathering and reporting
NOREPORT	Does not produce printed interval reports
NOTRACE	Does not trace any variables (no Trace Activity report)
NOVMGUEST	Does not measure CPU dispatch times and processor utilizations for a z/OS system running as z/VM guest.
OPTIONS	Prints a list of session options at the operator console at the start of the session, allowing the operator to change options. For a fast start-up of Monitor I, we recommend that you specify NOOPTIONS unless changes at start-up are really necessary.
STOP(8H)	Ends the session 8 hours after it was started
SYNC(SMF)	Synchronizes the reporting interval with SMF

Description of Monitor I data gatherer options

CACHE

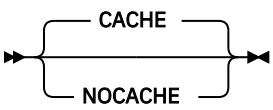

Specifies cache activity measurement. When you specify CACHE, or allow the default value to take effect, RMF gathers activity data for cache control units (there is no support for 3880 control units). The gathered data is stored in SMF records type 74 subtype 5.

Cache controller data is gathered by individual device address. There is no indication of which system in the sysplex initiates a recorded event. Therefore, the data can be gathered on any system sharing the cached devices. To avoid having duplicated data, you should gather cache activity data on one system

only. Refer to "Example that shows how to set up gathering options" in “Generalizing parmlib members” on page 20, which shows how to set up gathering options.

To suppress the gathering of cache data, specify NOCACHE.

RMF does no real-time reporting of cache activity, so if you wish to monitor this activity, the gatherer option RECORD must also be in effect for the session. The RECORD option takes effect by default.

CHAN

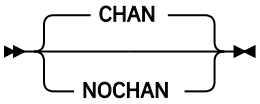

Specifies channel path activity measurement. A channel path is the physical interface that connects control units (grouped into logical control units) and devices to the channel subsystem.

CPU

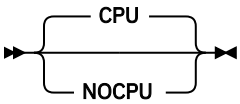

Specifies processor activity measurement.

CRYPTO

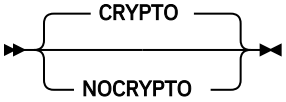

Specifies cryptographic hardware activity measurement.

CYCLE

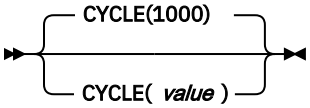

Specifies, in milliseconds, the length of the cycle at the end of which sampling observations are to be made, where **nnnn** is the number of milliseconds. The **valid range** is from a minimum of 50 to a maximum of 9999 milliseconds. If you specify less than 50, RMF will increase the value to 50. If you specify more than 9999, RMF will decrease the value to 9999. The **default value** is 1000 milliseconds. See “INTERVAL and CYCLE options” on page 68 for considerations that apply to choosing a cycle length.

DEVICE

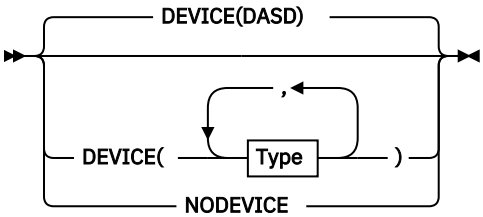

Type

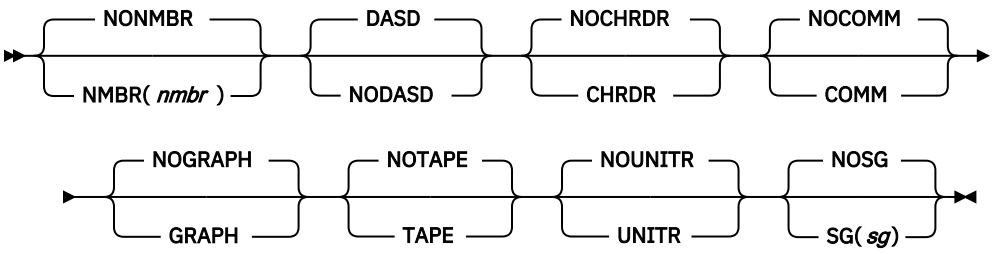

Specifies whether device activity is to be measured. You can request device activity by specifying all devices within one or more classes, or, optionally, one or more specific devices within each class. If you specify DEVICE, however, you must include an option; you need only include the classes you want to add to the default (DASD) or the specific device number you want data for.

Type is one of the following:

- One or more device numbers:

```
       ({[s]aaaa          })
{NMBR} ({[s]aaaa,[t]bbbb:[u]zzzz})/NONMBR
       ({[s]aaaa,[t]bbbb,... })
```

  NMBR requests specific device numbers, where aaaa, bbbb, and zzzz each represent hexadecimal 4-digit device numbers and s, t, and u each represent an optional 1-digit subchannel set ID. You can omit leading zeros. If the subchannel set ID is omitted data gathering for devices configured to any subchannel set is requested. You can specify any combination of:

  – a single device number, in the format [s]aaaa
  – a list of device numbers, in the format [s]aaaa,[t]bbbb
  – or a range of numbers in the format [t]bbbb:[u]zzzz, where [t]bbbb is your first number and [u]zzzz is your last number

  You can 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 [t]bbbb up to and including [u]zzzz.

  NONMBR, when specified, cancels any existing list of device numbers.

- 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

- One or more storage groups:

```
     ({aaaaaaaa                  })
{SG} ({aaaaaaaa,bbbbbbbb,...     })/NOSG
     ({aaaaaaaa,bbbbbbbb:zzzzzzzz})
```

  SG requests specific storage group names, where aaaaaaaa, bbbbbbbb, and zzzzzzzz each represent 1 to 8 character names. You can specify any combination of a single storage group name, in the format

aaaaaaaa, a list of names, in the format aaaaaaaa,bbbbbbbb,..., or a range of names, in the format bbbbbbbb:zzzzzzzz. Your entry can 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.

NOSG, when specified, cancels any existing lists of storage group names.

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.

Here are some examples of how to specify the DEVICE option.

Examples:

- To request device reporting for magnetic tape devices 0180, 0183, 0184, 0185, and 0188 as well as all direct access devices and communication equipment, you would specify:

```
DEVICE(COMM,NMBR(0180,0183:0185,0188))
```

  You do not need to specify DASD, because this is the default value.
- To request device reporting for magnetic tape devices and DASD you would specify:

```
DEVICE(TAPE)
```

  To limit the reporting of DASD, you must specify NODASD and use the NMBR field to identify those devices you want to monitor.
- If you request the following:

```
DEVICE(NODASD, NMBR(10288, 10291), 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.

ENQ

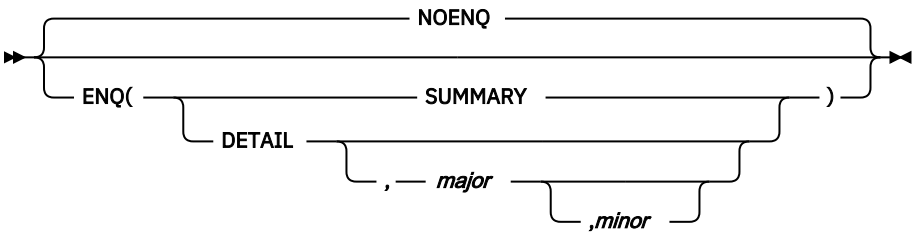

Specifies whether enqueue contention activity is to be measured. If you specify ENQ, you must specify either DETAIL or SUMMARY. When you specify DETAIL, the report includes the names of jobs that own the resource, have the longest period of contention, and are waiting for the resource. The names reported are selected during the period of maximum contention. When you specify SUMMARY, no names are reported. The default is NOENQ.

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 44 character minor name of the resource.

Example:

```
ENQ(DETAIL,SYSDSN,SYS1.PARMLIB)
```

To measure contention for a specific resource, use the name fields; to measure contention for all resources, do not specify a name. When you omit a minor name, all resources with the major name you specify are included.

ESS

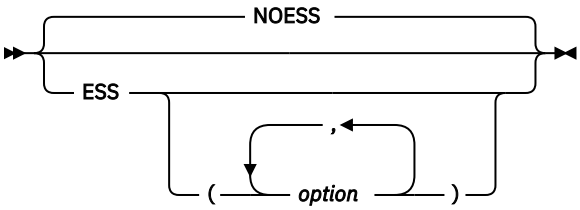

Specifies whether Enterprise Disk System statistics should be gathered. The gathered data is stored in SMF records type 74 subtype 8.

If you specify *option*, this can be one or more of the following:

LINK
asynchronous I/O as well as synchronous I/O link performance statistics are gathered.

RANK
extent pool statistics and rank statistics are gathered.

NOLINK
no link performance statistics are gathered.

NORANK
no extent pool and rank statistics are gathered.

If you do not specify neither the ***LINK*** nor the ***RANK*** option, then both ***LINK*** and ***RANK*** are default.

As ESS data gathering involves cache activity measurement (see option CACHE), it is recommended to specify both options in common. If you specify ESS together with NOCACHE, cache data is gathered implicitly without writing SMF 74 subtype 5 records.

In a sysplex, options CACHE and ESS can be specified on any system sharing the measured devices. Therefore specify options ESS and CACHE together on one selected system only to avoid duplicate data gathering.

Example:

To request all available statistics, specify one of the following:

- ESS
- ESS(LINK,RANK)

EXITS

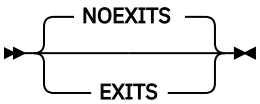

Specifies whether Monitor I user exit routines are to be executed during session processing to gather and report on additional data. See the *z/OS RMF Reporter Programmer's Guide* for information on using the exit routines.

If you have specified in the past the option EXITS to gather SMF records with the Cache RMF Reporter (CRR) Program Offering (records type 245), this is not required anymore with the Monitor I gathering option CACHE. Therefore, you should specify NOEXIT, unless you have some other exit routines that you want to activate.

FCD

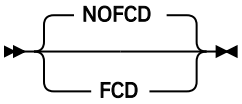

Specifies whether FICON director activities should be measured.

FICON director activity data is gathered by port address. There is no indication which system in the sysplex requested the I/O. Therefore, the data can be gathered on any system sharing the FICON directors. To avoid having duplicated data, you should set the FCD option on one system only.

Note: If you have specified the FCD option, please ensure that you do not disable the gathering of FICON director statistics on that system by setting STATS=NO in the IECIOSxx parmlib member. See the *z/OS MVS Initialization and Tuning Reference* for more information on the FICON STATS parameter.

INTERVAL

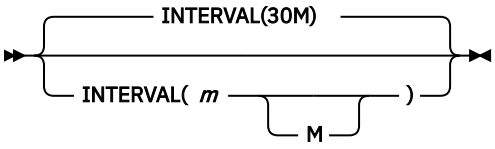

Specifies the length of the Monitor I reporting interval, where **m** is a divisor of 60, and **M** is minutes. This means that interval values of 1, 2, 3, 4, 5, 6, 10, 12, 15, 20, 30 or 60 minutes are possible, all of them meeting the SYNC value at least every hour.

At the end of the interval, the system dispatches Monitor I. Monitor I summarizes the recorded data and formats it into an interval report, or an SMF record, or both (see the Monitor I REPORT and RECORD options).

Note:

RMF processes this session option only if it is used in conjunction with one of the following **SYNC** options:

```
SYNC(RMF,mm)
SYNC(RMF,mmM)
NOSYNC
```

With SYNC(SMF), which is default, INTERVAL is ignored.

The default is 30 minutes (30M). The valid range for INTERVAL is from a minimum of one to a maximum of 60 minutes. If you specify less than one minute, RMF increases the value to one; if you specify more than 60 minutes, RMF decreases the value to 60. To synchronize the RMF reporting interval to any time within the hour, use the Monitor I SYNC option. See “Synchronizing SMF recording intervals” on page 17 for more information.

Note:

1. If you specify a STOP option, be sure that the **value** used there is equal to or greater than the INTERVAL value. Otherwise, RMF sets the STOP value to the INTERVAL value.
2. RMF extends INTERVAL in two situations:
   - When the system does not dispatch Monitor I at the end of the interval.

     If RMF is executing, and does not get control within the specified interval length, RMF will extend the length to 99 minutes. If RMF still does not get control within the 99 minutes, data collection is skipped for that interval, and RMF issues a message to the operator. This can happen when the dispatching priority for RMF is too low; see “Setting up the RMF control session including Monitor I and Monitor II” on page 15 on how to change the dispatching priority.
   - When you stop the processor during the interval.

If the processor is stopped during the interval, the interval length can also exceed 99 minutes. To avoid missing data collection, stop the RMF monitor or control session before stopping the processor.

IOQ

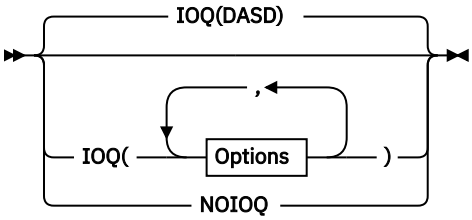

Options

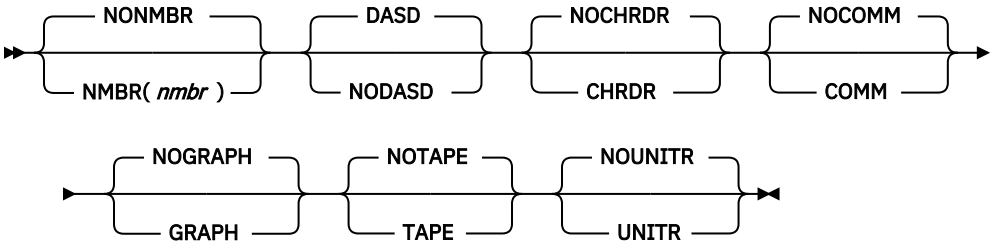

Specifies I/O queuing activity measurement for the devices in each logical control unit (LCU), where **option** can be any one of the following:

- One or more LCU numbers:

```
       ({aaaa          })
{NMBR} ({aaaa,bbbb:zzzz})/NONMBR
       ({aaaa,bbbb,... })
```

  NMBR requests specific logical control unit numbers, where aaaa, bbbb, and zzzz each represent hexadecimal logical control unit numbers. You can omit leading zeros. You can specify any combination of a single logical control unit number, in the format aaaa, a list of logical control unit numbers, in the format aaaa,bbbb, or a range of numbers in the format bbbb:zzzz, where bbbb is your first number and zzzz is your last number. You can 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 bbbb up to and including zzzz.

  NONMBR is the default; when specified, it cancels any existing lists of LCU numbers.

- 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. If you specify IOQ, you must include an option. The option need include only the classes you want to either add to the default (DASD) or the specific LCU number you want data for. The definition of an LCU is model-dependent.

On all processors, 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, but the I/O processor (SAP), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU of the 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.

Note: When your system is running as a guest under VM, RMF cannot gather data for it. In this case, the I/O Queuing Activity report shows only the static configuration data.

Example:

- To request I/O queuing activity for magnetic tape device LCUs 1130, 1133, 1134, 1135, and 1150 as well as all LCUs of the DASD and COMM classes, specify:

  ```
  IOQ(COMM,NMBR(1130,1133:1135,1150))
  ```

  LCUs of DASDs would be included by default, and the other device classes would be excluded by default.
- To request I/O queuing activity for LCUs for magnetic tape devices and DASD, specify:

  ```
  IOQ(TAPE)
  ```
- 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.

MEMBER

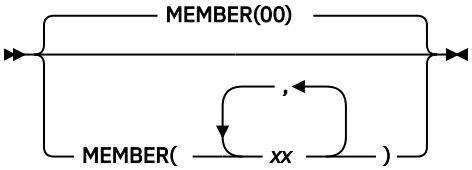

Specifies the Parmlib member(s) — up to five members can be specified — that contain Monitor I options for the session, where **(list)** contains from one to five members, separated by commas. Each member must be a two-character alphameric value. RMF then forms the member name by adding the two-character alphanumeric value to the ERBRMF prefix.

For the Monitor I session, the default is 00, indicating Parmlib member ERBRMF00. The contents of ERBRMF00 are described in “Storing gatherer options” on page 19. If you have created your own Parmlib data set, make sure you specify it in the RMF cataloged procedure. See “Setting up the RMF control session including Monitor I and Monitor II” on page 15.

If you specify more than one member, RMF processes the members' options in left to right priority order. For examples on how RMF processes session options, see Chapter 5, “How RMF processes session options,” on page 41.

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

Note: The MEMBER option must not be used in the RMF Parmlib members, but is only valid together with an RMF MODIFY console command.

OPTIONS

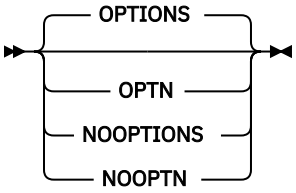

Specifies whether an options list for the session is to be printed at the operator console at the start of the session. If you specify OPTIONS, you can respond with any desired changes, except the MEMBER option, from the operator console.

To avoid unnecessary console output and delay in starting the session, specify NOOPTIONS. However, if RMF detects any syntax errors while processing session options, OPTIONS is forced.

Figure 3 on page 59 shows a console output with the Monitor I option OPTIONS in effect. For each option, this console output shows the source where the option has been set, for example, `-- COMMAND` means that the option has been set using a START or MODIFY command.

Table 5 on page 59 explains all possible sources which may appear in a console output.

```
ERB305I ZZ : PARAMETERS
ERB305I ZZ :   NOVMGUEST  -- DEFAULT
ERB305I ZZ :   WKLD  -- DEFAULT
ERB305I ZZ :   VSTOR(S)  -- DEFAULT
ERB305I ZZ :   NOTRACE  -- DEFAULT
ERB305I ZZ :   NOREPORT  -- DEFAULT
ERB305I ZZ :   SYSOUT(0)  -- MEMBER
ERB305I ZZ :   SYNC(SMF)  -- MEMBER
ERB305I ZZ :   NOSTOP  -- MEMBER
ERB305I ZZ :   RECORD  -- MEMBER
ERB305I ZZ :   PAGING  -- MEMBER
ERB305I ZZ :   PAGESP  -- MEMBER
ERB305I ZZ :   OPTIONS  -- MEMBER
ERB305I ZZ :   IOQ(NONMBR)  -- DEFAULT
ERB305I ZZ :   IOQ(UNITR)  -- MEMBER
ERB305I ZZ :   IOQ(TAPE)  -- MEMBER
ERB305I ZZ :   IOQ(GRAPH)  -- MEMBER
ERB305I ZZ :   IOQ(COMM)  -- MEMBER
ERB305I ZZ :   IOQ(CHRDR)  -- MEMBER
ERB305I ZZ :   IOQ(DASD)  -- MEMBER
ERB305I ZZ :   FCD  -- MEMBER
ERB305I ZZ :   EXITS  -- MEMBER
ERB305I ZZ :   ESS(RANK)  -- MEMBER
ERB305I ZZ :   ESS(LINK)  -- MEMBER
ERB305I ZZ :   ENQ(DETAIL)  -- MEMBER
ERB305I ZZ :   DEVICE(NOSG)  -- DEFAULT
ERB305I ZZ :   DEVICE(NONMBR)  -- DEFAULT
ERB305I ZZ :   DEVICE(UNITR)  -- MEMBER
ERB305I ZZ :   DEVICE(TAPE)  -- MEMBER
ERB305I ZZ :   DEVICE(GRAPH)  -- MEMBER
ERB305I ZZ :   DEVICE(COMM)  -- MEMBER
ERB305I ZZ :   DEVICE(CHRDR)  -- MEMBER
ERB305I ZZ :   DEVICE(DASD)  -- MEMBER
ERB305I ZZ :   CYCLE(1000)  -- MEMBER
ERB305I ZZ :   CRYPTO  -- MEMBER
ERB305I ZZ :   CPU  -- MEMBER
ERB305I ZZ :   CHAN  -- MEMBER
ERB305I ZZ :   CACHE  -- MEMBER
ERB305I ZZ :   MEMBER(10)  -- COMMAND
```

Figure 3. Console sample output with Monitor I OPTIONS in effect

Table 5. Where to specify Monitor I options

Source	Where Option is specified
-- COMMAND	On a START or MODIFY command.

Table 5. Where to specify Monitor I options (continued)	
Source	**Where Option is specified**
-- DEFAULT	In the program defaults.
-- EXEC	On the EXEC statement in the RMF cataloged procedure.
-- CHANGED	RMF changed the option. A message describes the conflict and the change RMF made.
-- MEMBER	In the RMF Parmlib member.
-- REPLY	The option was changed from the operator console in reply to message ERB306I.

PAGESP

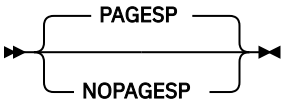

Specifies whether page data set activity is to be measured.

PAGING

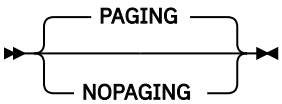

Specifies whether system paging activity is to be measured.

RECORD

Specifies whether measured data is to be written to SMF records. In order for RECORD to take effect, the complementary SMF enabling procedures must first be performed. These procedures are described in *z/OS MVS System Management Facilities (SMF)*.

Note: If you specify NORECORD, do not specify the NOREPORT option at the same time. RMF changes NOREPORT to REPORT(DEFER) if you do.

REPORT

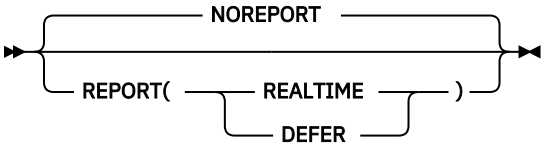

Specifies whether printed interval reports of the measured data are to be produced. This option is ignored for the Workload Activity report if the system is running in goal mode. Request this report from the Postprocessor, using the SYSRPTS option. When you omit the option, the default is NOREPORT. If you specify REPORT, you must specify either REALTIME or DEFER.

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.

Note:

1. If you specify NOREPORT, do not specify the NORECORD option at the same time. RMF changes NOREPORT to REPORT(DEFER) if you do.
2. If you specify REPORT(DEFER), do not specify the NOSTOP option at the same time. If you do, RMF changes NOSTOP to STOP with a value equal to the INTERVAL value.

STOP

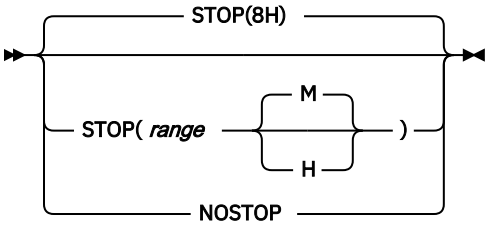

Specifies the desired duration for the Monitor I session in minutes (M) or hours (H). The **valid range** is from a minimum of one minute to a maximum of one week (168 hours or 10,080 minutes). If you do not specify a value, the **default range** is 8 hours. If you specify less than one minute, RMF will increase the value to one minute; if you specify more than 168 hours, RMF will decrease the value to 168 hours. 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

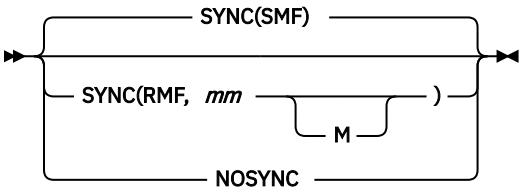

Specifies whether the interval is to be synchronized with SMF, or on the minute with the RMF interval synchronization mechanism.

SYNC(SMF) is the default and specifies that RMF will synchronize its interval using SMF's global interval and synchronization values.

The **valid range** is the number of minutes from 0 to 59 (mm), past the hour at which synchronization is to occur. If any value other than 0 through 59 is specified, or the value is omitted, 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 is a factor of 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 specifies that no synchronization is to be performed. Do not specify this if you want to generate sysplex reports.

Note:

1. If you specify SYNC(SMF), do not specify the INTERVAL option at the same time. If you do, RMF ignores the INTERVAL specification.
2. If you use the syntax for the SYNC option from a release prior to RMF 4.3.0, that is, SYNC(nn), this will automatically be converted to SYNC(RMF,nn).

See "Synchronizing SMF recording intervals" on page 17 for more information.

SYSOUT

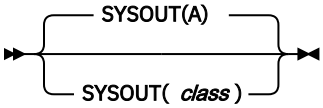

Specifies the SYSOUT class to which the formatted interval reports are directed. Class A is the default. The SYSOUT option cannot be modified during the session.

TRACE

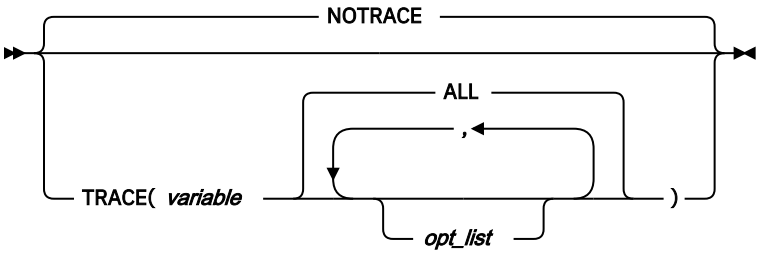

Specifies whether to trace certain variables for the Trace Activity report.

Note: Monitor I gathers and reports all trace variables as they are provided by the system. It has no influence on the type and does not perform any calculation.

Valid variables are:

Variable
Value

ASMERRS
bad slots on local page data sets

ASMIORQC
count of I/O requests completed and returned to RSM

ASMIORQR
count of I/O requests received by I/O control

ASMNVSC
total local slots allocated for non-VIO private area pages

ASMSLOTS
total local slots (sum of slots in open local page data sets)

ASMVSC
total local slots allocated for VIO private area pages

CCVCPUCT
number of online CPUs

CCVEJST
this variable is no longer supported

CCVENQCT
number of users non-swappable for enqueue reasons

CCVRBSTD
recent base time of day

CCVRBSWT
recent base system wait time

CCVUTILP
system CPU utilization

LSCTCNT
current number of logically swapped users for terminal wait

LSCTMTE
maximum think time allowed for logical swap candidate

MCVFRCNT
number of pages needed to be stolen by force steal routine

MCVMGAGE
expanded storage migration age

MCVSBLTF
long term percentage of eligible storage that is actually fixed

MCVSIPR
common page-in rate

MCVSTCRI
highest system UIC

MCVTWSS
common target working set size

OMDGAMRE
maximum number of messages on the action message retention facility (AMRF) queue. If a large number of action messages are retained on the AMRF queue for a particular period, it may mean more operators are needed for that period.

OMDGCMDI
number of commands issued per second.

OMDGOREB
maximum number of operator reply entries (OREs) on the system reply queue. To eliminate thrashing, use this number to monitor and adjust the ORE buffer limit set at IPL time. To dynamically adjust this limit, use the CONTROL M command.

OMDGWQEB
maximum number of WTO queue elements (WQEs) on the system output queue. To eliminate thrashing (excessive data movement which confines system to doing little useful work), use this number to monitor and adjust the WTO buffer time limit set at IPL time. To dynamically adjust this limit, use the CONTROL M command.

OMDGWTLI
number of write-to-logs (WTLs) issued per second, indicating the number of records going to SYSLOG within a time period. To control the number of data sets produced during the day, vary the number of records per SYSLOG data set.

OMDGWTOI
total number of lines of messages, write-to-operators (WTOs) issued per second. Use it to determine the peak message rate period and the average message rate.

RAXESCT
number of common storage pages on expanded storage

RAXFMCT
number of frames allocated to common storage

RCEAEC
total number of expanded storage E frames currently on the ESTE queue

RCEAECLO
available expanded storage low threshold

RCEAECOK
available expanded storage satisfactory threshold

RCEAFC
total number of frames currently on all available frame queues

RCEAFCLO
available central storage low threshold

RCEAFCOK
available central storage satisfactory threshold

RCEBELFX
total number of fixed pages below 16 megabytes in central storage, which is the sum of page-fixed LSQA, SQA (excluding reserved SQA) and V=R allocated pages.

RCECOMPI
number of common area pages paged-in

RCECOMPO
number of common area pages paged-out

RCEDFRS
number of times a deferred frame allocation has been satisfied

RCEESINU
number of in-use expanded storage frames

RCEESREA
number of non-VIO pages read from expanded storage

RCEESWRT
number of pages written to expanded storage frames

RCEHSPEM
total number of hiperspace pages migrated from expanded storage to auxiliary storage

RCEHSPER
total number of hiperspace pages in the system read from expanded storage to central storage

RCEHSPEW
total number of hiperspace pages written from central storage to expanded storage

RCEHSPPI
total number of hiperspace pages paged in from auxiliary storage

RCEHSPPO
total number of hiperspace pages paged out to auxiliary storage

RCELPAPI
number of PLPA and PLPA directory pages paged-in

RCEMVBEL
number of pages moved from below 16 megabytes in central storage

RCENWSF
total number of secondary and non-working set pages migrated to auxiliary storage.

RCEPAGMV
number of times a frame was moved from one frame to another

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

RCESPFR
number of frames available by swap-out without requiring I/O

RCESWPPI
total number of pages requiring I/O to swap-in

RCESWPPO
total number of pages requiring I/O to swap-out

RCETOTFX
total number of pages currently fixed, the sum of page-fixed LSQA, SQA (excluding reserved SQA) and V=R allocated pages

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

RCEVIOME
number of VIO pages written to expanded storage

RCEVIOMG
number of VIO pages migrated from expanded 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

RCEVIORE
number of VIO reads from extended storage

RCEWSDNE
total number of primary working set pages migrated to auxiliary storage

RCVAFQA
average available frame count

RCVAVQC
AVQ low count

RCVCPUA
CPU usage average * 16

RCVFXIOP
percentage of central storage that is fixed or allocated for paging

RCVMFXA
average number of fixed frames for the system

RCVPAGRT
total paging rate

RCVPTR
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

RCVUICA
UIC average

RMCAAWSC
APPC/MVS transaction scheduler (ASCH) wait swap count

RMCADWSC
detected wait physical swap count

RMCAEXSC
exchange on recommendation value swap count

RMCAFHLD
number of swaps failed because of an outstanding HOLD SYSEVENT

RMCAICSC
improve central storage use

RMCAIPSC
improve system paging rate

RMCALWSC
long wait physical swap count

RMCAMRSC
make room to swap in a user who was swapped out too long.

RMCANQSC
CPU enqueue exchange swap count

RMCAOISC
OMVS input wait

RMCAOOSC
OMVS output wait

RMCARSSC
central storage shortage swap count

RMCATISC
terminal input swap count

RMCATOSC
terminal output swap count

RMCATSSC
count of transition swaps

RMCAUSSC
unilateral swap out count

RMCAXSSC
auxiliary storage shortage swap count

RMCTTRPC
number of pages used for transaction elements

SMCABFLS
number of records lost because of a shortage of buffers

SMCABFWT
number of buffers written

SMCADSCT
number of records lost because of a full data set

SMCANMFL
current number of full buffers

SMCARCWT
number of records written

You can specify one or more of the following for **opt_list**:

MIN
minimum sampled value of the variable over the sampling period

MAX
maximum sampled value of the variable over the sampling period

AVG
average value of the variable over the sampling period

END
snapshot of the last value in the sampling period

STDDEV
standard deviation from the values sampled

ALL
default for **opt_list**, meaning all of the above

VSTOR

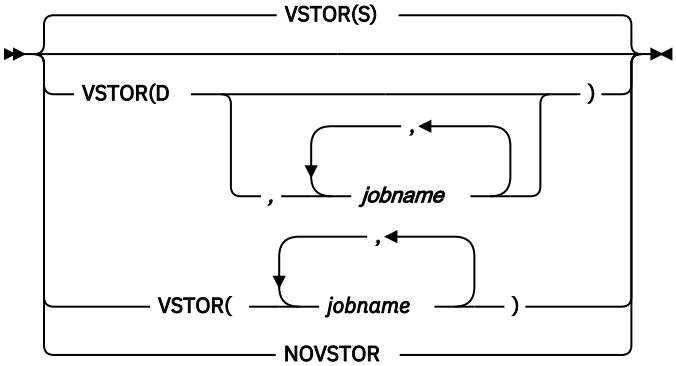

Specifies whether virtual storage activity is to be measured. RMF can produce common storage summary and detail reports and private area summary and detail reports. When you specify S, either explicitly or by default, RMF produces summary reports; when you specify D, RMF produces both summary reports 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 efficient to minimize the time required to gather the data by specifying one or two jobs separately. When selecting specific jobs, note also that RMF can gather meaningful data only for long-running jobs.

Note: Measuring virtual storage activity for a specific job may have significant impact on the performance of the job. System address spaces like CATALOG, VTAM®, DB2®, IMS or other, should be specified as *jobname* only for a short period of time when diagnosing a special performance situation. For VSTOR data gathering considerations, refer to the VSTOR report description in *z/OS RMF Report Analysis*.

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. Some other possible combinations are:

Examples:

- 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.

Note: Modifications on the VSTOR option are always treated as add-on. For example, when the current status is VSTOR(D,jobname1) and you specify VSTOR(jobname2) via the Modify command, the result will be VSTOR(D,jobname1,jobname2). Now, when you specify VSTOR(S) via a Modify, the status of the VSTOR option is not changed at all because S (summary) is already part of D (detail). VSTOR(D) tells you that Summary and Detail are active. Resetting a VSTOR parameter is only possible by specifying NOVSTOR followed by another VSTOR(...).

WKLD

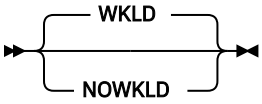

With the WKLD option, you specify whether the system workload is to be measured. WKLD is the default, so measuring will be done automatically, unless you specify NOWKLD.

VMGUEST

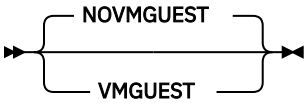

With the VMGUEST option, you specify whether CPU dispatch times and processor utilizations should be measured for a z/OS system when this system is running as z/VM guest. In such a case, you can request a simplified Partition Data Report.

Special considerations

Specify Monitor I session options carefully. If RMF detects any conflicting options when processing session options, it selects compatible values for them, and reports the changes in a message to the operator console.

Other groups of options do not cause actual conflicts, but you must choose their values carefully to avoid undesirable results. These options include:

- INTERVAL and CYCLE options
- STOP, INTERVAL, and REPORT options
- Device class selection for the DEVICE option

INTERVAL and CYCLE options

Much of the data in the Paging, Page Data Set, Processor, Trace, Virtual Storage, CPU, Channel, I/O queuing, and Device Activity reports is statistically sampled. As the accuracy of sampled data increases with the number of random samples taken, 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.

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.

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 reduce RMF overhead further, select specific devices for reporting rather than entire device classes. In the case of Postprocessor routines, selecting specific devices can result in shorter reports, thus saving both time and paper. 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. Therefore, the cycle value should not be made too small, especially when the number of UCBs for measured device classes is large.

Chapter 7. Snapshot data gathering with Monitor II

You can run Monitor II as background session to create SMF type 79 records.

This session is started by the operator, and all options are defined in Parmlib member ERBRMF01 or by operator commands.

All valid options are similar to those you can use during a Monitor II display session, so they are described in Snapshot reporting with Monitor II in *z/OS RMF Reporter User's Guide*.

Chapter 8. Short-term data gathering with Monitor III

This information unit describes:

- the syntax and effect of the Monitor III data gathering options
- how to control VSAM data set recording
- how Monitor III data gathering handles the daylight saving time

The detailed descriptions of the options are in alphabetical order.

Summary of gatherer session options

You can specify Monitor III gatherer session options before or during the session.

Before the session, use the following:

- The Monitor III gatherer session parmlib member. The default member is ERBRMF04. See "Storing gatherer options" on page 19 for its contents. For a description of the MEMBER option and how to specify other parmlib members, see "Description of Monitor III data gatherer options" on page 75.
- The **parm** field of the START session command that starts the session. See "Starting a specific Monitor" on page 35.

During the session, use the following:

- The **parm** field of the MODIFY session command, to modify options already in effect. See "Modifying RMF session options" on page 37.
- The response to the OPTIONS option.

Table 6 on page 73 gives a summary of the Monitor III gatherer session options. The referenced pages describe the options in detail.

Table 6. Monitor III Data Gatherer Session Options

Option	Effect	Details on
CACHE(suboption...)	Defines cache data gathering.	"CACHE" on page 75
CFDETAIL	Defines level of detail for data gathering for the coupling facility.	"CFDETAIL" on page 75
CRYPTO	Controls data gathering for cryptographic hardware activity reports.	"CRYPTO" on page 76
CYCLE(nnnn)	Sets the length of the cycle at the end of which RMF samples data.	"CYCLE" on page 76
DATASET(suboption...)	Controls data set recording of sampled data.	"DATASET" on page 76
EADM	Controls data gathering for extended asynchronous data mover (EADM) activity report.	"EADM " on page 77
HFSNAME(suboption...)	Controls data set recording for z/OS UNIX file systems.	"HFSNAME" on page 77
IOSUB	Controls data set recording of I/O-subsystem and channel-path activity.	"IOSUB" on page 77
LOCK	Defines data gathering for lock reporting (spin locks and suspend locks).	"LOCK" on page 78
MASTER	Makes the system eligible/uneligible to be the RMF Master Gatherer system.	"MASTER" on page 78
MEMBER(list)	Specifies Parmlib members containing session options.	"MEMBER" on page 78
MINTIME(nnn)	Specifies the interval at which data samples are summarized.	"MINTIME" on page 79
OPD	Defines data gathering for OMVS process data.	"OPD" on page 79
OPTIONS	Controls display of the current options at the start of a session.	"OPTIONS" on page 79

Table 6. Monitor III Data Gatherer Session Options (continued)		
Option	**Effect**	**Details on**
PCIE	Controls data gathering for PCIE activity report.	“PCIE” on page 80
RESOURCE(...)	Specifies the job entry subsystem (JES) to be used.	“RESOURCE” on page 80
SGSPACE(suboption...)	Defines data gathering for storage group space and disk space monitoring.	“SGSPACE” on page 81
SCM	Deprecated. Use the EADM option instead.	“EADM ” on page 77
STOP(value)	Sets the duration of the data gatherer interval.	“STOP” on page 81
SYNC	Synchronizes MINTIME within the hour.	“SYNC” on page 82
SYSOUT(class)	Specifies the SYSOUT class for gatherer messages.	“SYSOUT” on page 82
VSAMRLS(suboption...)	Controls data gathering for VSAM RLS activity.	“SYSOUT” on page 82
WSTOR	Sets the size of the RMF local storage buffer.	“WSTOR” on page 83
ZFS	Defines data gathering for monitoring zFS activity.	“ZFS” on page 83
ZIIPUSE	Specifies whether the Monitor III data gatherer is entitled to execute partially on IBM Z Integrated Information Processors (zIIPs).	“ZIIPUSE” on page 83

Default gatherer session options

Here are the options that take effect by default. You need to specify an option only if you want to change that option to a different value:

Table 7. Monitor III Default Session Options	
Default Option	**Description**
CACHE	Defines cache data gathering.
CFDETAIL	Defines partial data gathering for the coupling facility.
CRYPTO	Activity data is gathered for cryptographic hardware.
CYCLE(1000)	Takes data samples once a second (1000 milliseconds).
DATASET(STOP,NOSWITCH)	No data set recording will be done.
EADM	Activity data is gathered for extended asynchronous data mover (EADM).
IOSUB	Defines data gathering for the I/O subsystem and for channels.
MASTER	Makes the system eligible to be the RMF Master Gatherer system.
MINTIME(100)	Builds a set of samples every 100 seconds.
NOLOCK	No data gathering for lock reporting (spin locks and suspend locks).
NOOPTIONS	Session options are not displayed at the operator console at the start of the session.
NOSGSPACE	No data gathering for storage group and disk space monitoring.
NOSTOP	The session does not stop automatically after a predefined time; you must use a STOP command.
OPD	Defines data gathering for OMVS process data.
PCIE	Activity data is gathered for PCI Express based functions.
RESOURCE(*JES2,JES2)	Assumes that JES2 is installed on the system.
SYNC(0M)	MINTIME is synchronized on the hour.
VSAMRLS	Activity data is gathered for VSAM RLS by storage class.
WSTOR(128)	Sets the RMF local storage buffer to 128 megabytes.
ZFS	Activity data is gathered about zFS.

Table 7. Monitor III Default Session Options (continued)	
Default Option	**Description**
ZIIPUSE	Specifies whether the Monitor III data gatherer is entitled to execute partially on IBM Z Integrated Information Processors (zIIPs).

Monitor III creates two types of records:

Set of samples
These records are written into the local storage buffer and (if specified on the DATASET option) into VSAM data sets.

SMF records
These records are written if defined in the SMFPRMxx Parmlib member.

You can find detailed information about all record types in Table 2 on page 5.

Description of Monitor III data gatherer options

CACHE

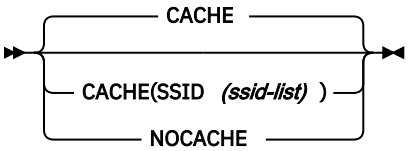

Specifies cache activity measurement. When you specify CACHE, or allow the default value to take effect, RMF gathers activity data for cache control units (there is no support for 3880 control units).

Cache controller data is gathered by individual device address. There is no indication of which system in the sysplex initiates a recorded event. Therefore, the data can be gathered on any system sharing the cached devices.

Note: To avoid unnecessary high CPU utilization and duplicated data, you should gather cache activity data on one system only. Refer to "Example that shows how to set up gathering options" in “Generalizing parmlib members” on page 20, which shows how to set up gathering options.

To suppress the gathering of cache data, specify NOCACHE.

CFDETAIL

Controls the collection of data about the coupling facility. If this option is active, detail data about activities in the structures (LIST, LOCK, and CACHE) of the coupling facility will be stored in the set-of-samples area, and can be seen in the Coupling Facility Activity report.

This data collection is optional. The default is CFDETAIL. To prevent detailed data collection, specify NOCFDETAIL when starting the Monitor III session or use a MODIFY command during a running session. Specifying NOCFDETAIL on a MODIFY command stops the data collection at the end of the current Mintime.

With CFDETAIL, additional data is being gathered that enables you to get many details about the usage of each structure in the coupling facility. Consider that this data gathering is done only on the RMF Master Gatherer system (see also “MASTER” on page 78).

CRYPTO

Controls the collection of activity data for cryptographic hardware. The default value is CRYPTO.

CYCLE

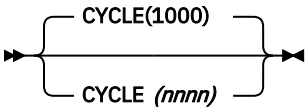

Specifies the length of a cycle at the end of which RMF samples data, where *nnnn* is the length in milliseconds. The valid range value is 50 to 9999. If you specify a value outside the valid range, RMF uses 9999 milliseconds for values above the range and 50 milliseconds for values below it.

Sysplex Reporting:

Please use the same CYCLE value for all systems in the sysplex to enable correct sysplex reporting.

The default value is 1000 milliseconds (one second). Decreasing the CYCLE value to less than one second brings little improvement in the quality of the statistics produced, compared to the following adverse effects on performance:

- Increasing the amount of processor time needed to sample data
- Causing RMF to fill the wrap-around storage buffer more quickly
- Using more space in the user-defined VSAM data set

DATASET

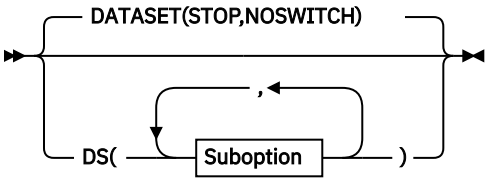

Suboption

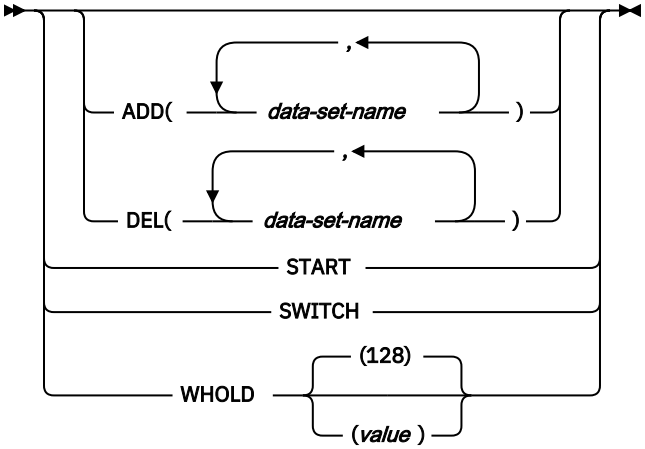

Controls the recording of samples in user-defined data sets. The suboptions are:

- ADD|DEL
- START|STOP

- SWITCH|NOSWITCH
- WHOLD

For detailed information on the DATASET option and its suboptions, see "Controlling data set recording" on page 84.

EADM

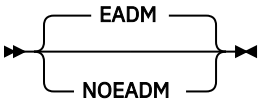

Controls the collection of activity data for extended asynchronous data mover (EADM). The default is EADM.

Note: This gatherer option was formerly called SCM/NOSCM. The SCM and NOSCM keywords are accepted with the same meaning as EADM/NOEADM.

HFSNAME

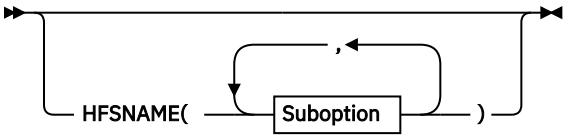

Suboption

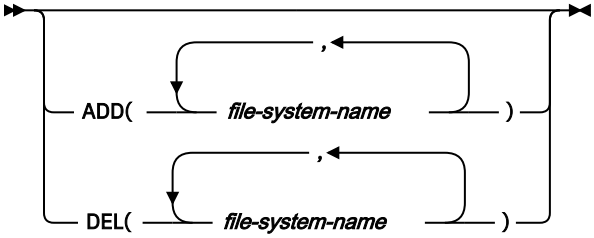

Controls the recording of statistics for UNIX HFS names. The suboptions are:

- ADD - Start data gathering for a UNIX hierarchical file system (HFS)
- DEL - Stop data gathering for a UNIX hierarchical file system

This data gathering is required to create the File System Statistics part of the HFS Postprocessor report.

IOSUB

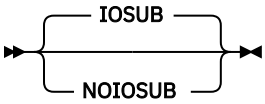

Controls the collection of data about the I/O subsystem configuration. I/O-queuing and channel-path activities can be stored in the set-of-samples area.

This data collection is optional. The default is IOSUB. To stop collection, specify NOIOSUB when starting or modifying the Monitor III session. When you specify IOSUB on a MODIFY command, collection starts at the end of the current Mintime.

LOCK

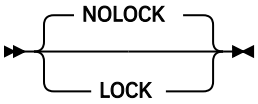

Controls data gathering about spin locks and suspend locks. The default is no data gathering.

MASTER

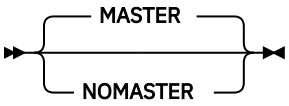

One member of a sysplex is selected by RMF to gather Monitor III data. This is called *sysplex master gathering* and has been implemented to reduce workload on non-master members and to reduce the amount of data in SSHs and SMF records. The RMF Master Gatherer system is determined by the RMF Sysplex Data Server automatically according to the following set of rules:

1. Monitor III gatherer active
2. highest z/OS release
3. Sysplex Data Server running with SMF buffer (SMFBUF option)
4. MASTER option specified

You can use the MASTER parmlib option to refine the determination, which of the systems in a sysplex becomes the RMF Master Gatherer system. If the MASTER option is specified for a certain system, this system is one candidate for the MASTER status. If NOMASTER is set, it will not be a MASTER candidate if there are other eligible systems having the MASTER option set. Thus you can use a combination of MASTER/NOMASTER options to select the RMF Master Gatherer system, if there are several systems, that fulfil the priorities “1” on page 78 through “3” on page 78 rules simultaneously.

Beyond using the MASTER option, the Monitor III Master Gatherer status of a system can be changed dynamically by means of MODIFY commands as described in “Modifying RMF session options” on page 37.

MEMBER

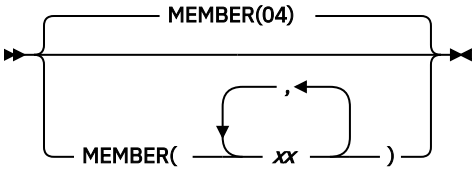

Specifies one to five Parmlib members that contain Monitor III gatherer options for the session. Each member is represented by a two-character alphameric value, to which RMF adds to the prefix ERBRMF to form the member name. The values in **(list)** must be separated by commas.

For the Monitor III gatherer session, the default is 04, indicating Parmlib member ERBRMF04. If you have created your own Parmlib, make sure you specify it on the IEFRDER DD statement in the RMF cataloged procedure. See “Setting up the RMF control session including Monitor I and Monitor II” on page 15.

If you specify an option in more than one member, RMF uses the value specified in the leftmost member of the list.

MINTIME

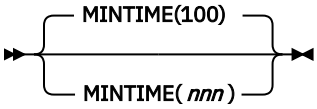

Specifies, in seconds, the length of a time interval. At the end of this interval, the data gatherer combines all samples it has gathered into a set of samples. The samples combined at the end of each MINTIME interval can then be summarized and reported by the data reporter.

Sysplex Reporting:

Use the same MINTIME value for all systems in the sysplex to enable correct sysplex reporting.

Valid MINTIME values range from 10 to 999. The default is 100. If you specify a value outside the valid range (10 to 999), RMF uses 999 seconds for values above the range and 10 seconds for values below the range. MINTIME is the smallest time interval the data reporter can report on.

See "Synchronizing SMF recording intervals" on page 17 for more information about using MINTIME values to synchronize Monitor I and Monitor III recording intervals.

OPD

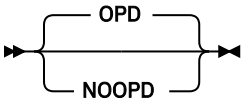

Specifies measurements for OMVS process data.

OPTIONS

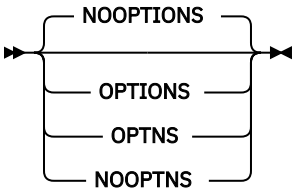

Specifies whether or not an option list for the session is to be printed at the operator console at the start of the session. If you specify OPTIONS, the list is printed, and you can respond with any desired changes, except to the MEMBER option, from the operator console.

If you do not want to make any changes, you should specify NOOPTIONS. This saves time when starting the session. However, if RMF detects any syntax errors while processing session options, OPTIONS is forced.

Figure 4 on page 80 shows the console output produced when OPTIONS is in effect and seven data sets are specified for data set recording. (See "Controlling data set recording" on page 84.)

The keywords on the right in the console output indicate from which source the current value for each option was taken. The meanings of the keywords are:

Table 8. Monitor III OPTIONS Command Sources

Keyword	**Source from which option was taken**
COMMAND	A START or MODIFY command.
DEFAULT	The program defaults.
EXEC	The EXEC statement in the RMF cataloged procedure.

Table 8. Monitor III OPTIONS Command Sources (continued)	
Keyword	**Source from which option was taken**
CHANGED	RMF changed a conflicting option. A message describes the conflict and the change RMF made.
MEMBER	The RMF Parmlib member.
REPLY	From the operator console in reply to message ERB306I.

```
ERB305I  III : PARAMETERS
ERB305I  III :  CACHE -- DEFAULT
ERB305I  III :  CFDETAIL -- DEFAULT
ERB305I  III :  CYCLE (1000) -- DEFAULT
ERB305I  III :  DATASET(STOP) -- DEFAULT
ERB305I  III :  DATASET(SWITCH) -- COMMAND
ERB305I  III :  DATASET(WHOLD(128)) -- DEFAULT
ERB305I  III :  DATASET(ADD(any.ds.name1)) -- MEMBER
ERB305I  III :  DATASET(ADD(any.ds.name2)) -- MEMBER
ERB305I  III :  DATASET(ADD(any.ds.name3)) -- MEMBER
ERB305I  III :  DATASET(ADD(any.ds.name4)) -- MEMBER
ERB305I  III :  DATASET(ADD(any.ds.name5)) -- MEMBER
ERB305I  III :  DATASET(ADD(any.ds.name6)) -- MEMBER
ERB305I  III :  DATASET(ADD(any.ds.name7)) -- MEMBER
ERB305I  III :  DATASET(WHOLD(128) -- DEFAULT
ERB305I  III :  WSTOR(128) -- DEFAULT
ERB305I  III :  MINTIME (100) -- DEFAULT
ERB305I  III :  NOSTOP -- DEFAULT
ERB305I  III :  SYNC(0) -- DEFAULT
ERB305I  III :  IOSUB -- DEFAULT
ERB305I  III :  OPD -- DEFAULT
ERB305I  III :  VSAMRLS -- DEFAULT
ERB305I  III :  OPTIONS -- COMMAND
ERB305I  III :  RESOURCE(*JES2,JES2) -- MEMBER
ERB305I  III :  SYSOUT(A) -- DEFAULT
ERB305I  III :  MEMBER (04) -- COMMAND
ERB305I  III :  NOSGSPACE -- DEFAULT
ERB305I  III :  ZFS -- DEFAULT
ERB305I  III :  PCIE -- DEFAULT
ERB305I  III :  EADM -- DEFAULT
ERB305I  III :  CRYPTO -- DEFAULT
```

Figure 4. Console Output with OPTIONS in Effect

PCIE

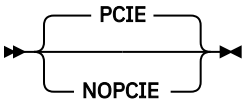

Controls the collection of activity data for PCI Express based functions. The default is PCIE.

RESOURCE

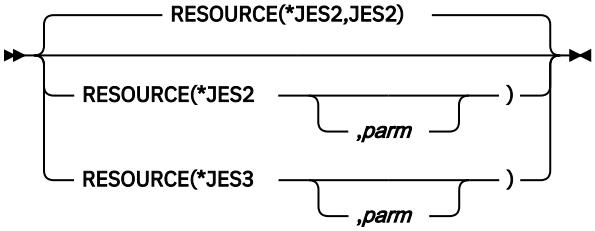

Specifies the job entry subsystem (JES) resource from which an address space requests service.

***JES2**
Required if the installed primary JES is JES2.

***JES3**
Required if the installed primary JES is JES3.

parm
This is an optional parameter. If your installation has chosen a name other than JES2 or JES3, then you must specify that name under **parm**.

The default is RESOURCE(*JES2,JES2).

SCM

Deprecated. Use the EADM option instead. See “EADM ” on page 77.

SGSPACE

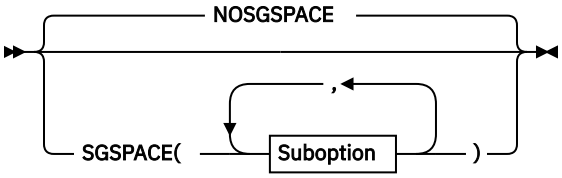

Suboption

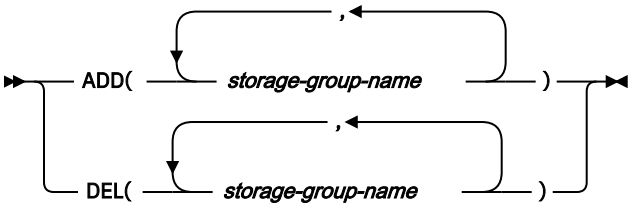

Controls data gathering for storage group space and disk space monitoring:

- You may specify multiple ADD/DEL suboptions.
- A storage group name must not be longer than 30 characters, otherwise it is ignored.
- You can specify up to 25 storage group names. Additional names are ignored.

Note: In a sysplex environment, it is recommended to activate the SGSPACE option for a certain storage-group-name on one system only to avoid duplicate data.

STOP

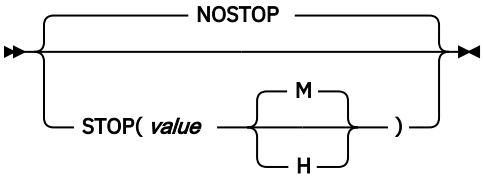

Specifies the desired duration for the data gatherer interval in minutes (M) or hours (H). You can specify a value from one minute (1M) to one week (168H or 10080M). RMF uses 168H for values above the range. If you do not specify M or H, RMF uses minutes (M).

NOSTOP means that only the session or system command STOP can end the session.

Note: The STOP option applies only to the data gatherer. The operator can use the session command STOP to end the session at any time, regardless of the value specified for this option. The RMF control session remains active until the operator issues a system command STOP.

SYNC

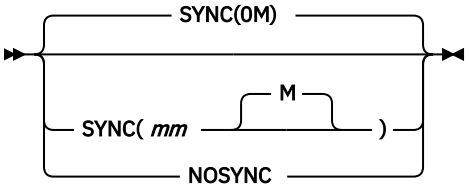

Specifies how the MINTIME interval is to be synchronized with the hour. This option must be specified if you want to generate sysplex reports. See “Synchronizing SMF recording intervals” on page 17 for more information. If you want synchronization, specify SYNC and the number of minutes (mm) after the hour (in a range from 0 to 59) at which you want synchronization. If you specify a value that is not between 0 and 59, RMF uses 0, the default, which synchronizes sets of samples on the hour. If you specify NOSYNC, all intervals are the same.

Note: Keep in mind the time you start a Monitor III data gatherer session. RMF synchronizes the starting time of a set of samples by calculating how many sets of samples will fit in the time range up to the first synchronization point. This might mean that the MINTIME interval before the synchronization point is shortened. Subsequent sets of samples remain synchronized only when the MINTIME value is a factor of 60.

SYSOUT

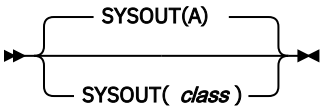

Specifies the SYSOUT class for messages generated by the data gatherer. You cannot modify the SYSOUT option while the data gatherer is active.

The default value is A.

VSAMRLS

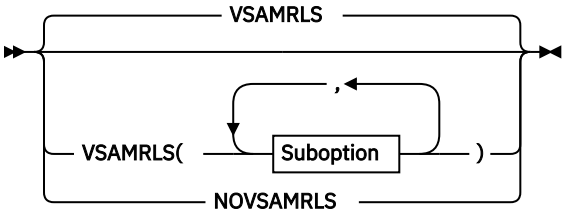

Suboption

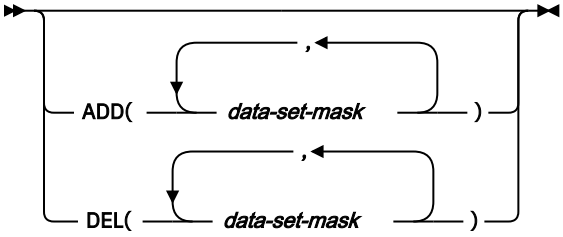

This option controls the collection of VSAM RLS activity data. By default, or if you specify VSAMRLS, activity data is gathered for VSAM RLS by storage class. In addition, you can specify data set masks to collect data by VSAM spheres. To suppress the gathering of VSAM RLS data, specify NOVSAMRLS.

You can control the collection of VSAM RLS activity data by VSAM spheres using following suboptions:

- ADD - Start collection for all VSAM data sets which are covered by the mask.

- DEL - Stop collection for all VSAM data sets which are covered by the mask.

Up to 50 different data set masks can be active at a time. You can not add a set of data sets by using the wildcard sign and afterwards delete a subset which is covered by the mask. For example, if VSM1.* has been added, you can not delete VSM1.VSR1.*.

A data set mask must apply to following rules:

- The data set mask represents a base cluster name. All components belonging to the base cluster will be gathered (data, index, alternate data, alternate index).
- The data set mask can be a full or partial data set name. For example, VSM1.VSR1.BASE or VSM1.*
- At least a high level qualifier must be specified.
- * specifies one qualifier, ** specifies any number of qualifiers
- Once a wildcard is specified, then no other qualifiers are allowed

Note: Since VSAM RLS Activity by VSAM spheres is a sysplex-wide report, the same set of data set masks should be active on all systems in the sysplex.

WSTOR

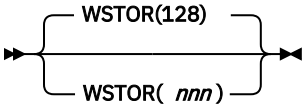

Specifies, in megabytes, the maximum size of RMF's local storage buffer for the data gatherer. The size of buffer that the data gatherer gets is either the value specified in this option or the maximum GETMAIN size defined by the system, whichever is smaller.

The valid range value is 4 to 999. RMF uses a default of 128 if you do not specify a value. If you specify a value outside the valid range, RMF uses 999 megabytes for a value above the range and 4 megabytes for a value below the range.

RMF stores the set of samples collected during a MINTIME in its own local storage buffer. If you specify data set recording during a session, RMF copies each set of samples from the local storage buffer to the currently active data set for the session. Common data items for a set of samples (such as jobname or device name) are held in tables to reduce the amount of local storage needed.

Note:

1. This option cannot be modified by the session command MODIFY.
2. When you specify the **WSTOR** parameter, you must ensure that there is enough space on the page data set to accommodate a buffer of the specified size.

ZFS

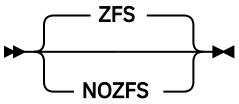

Specifies whether data gathering should (ZFS) or should not (NOZFS) be done for zFS file system activity. The default value is ZFS.

ZIIPUSE

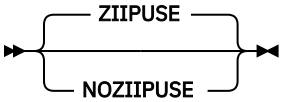

Specifies whether the Monitor III data gatherer is entitled to execute partially on IBM Z Integrated Information Processors (zIIPs).

Controlling data set recording

You control the recording of samples to the VSAM data sets through the data gatherer option DATASET. The syntax is:

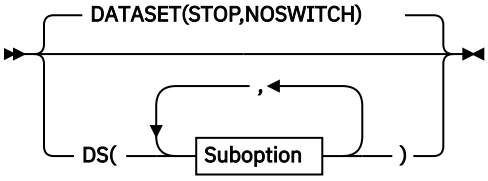

Suboption

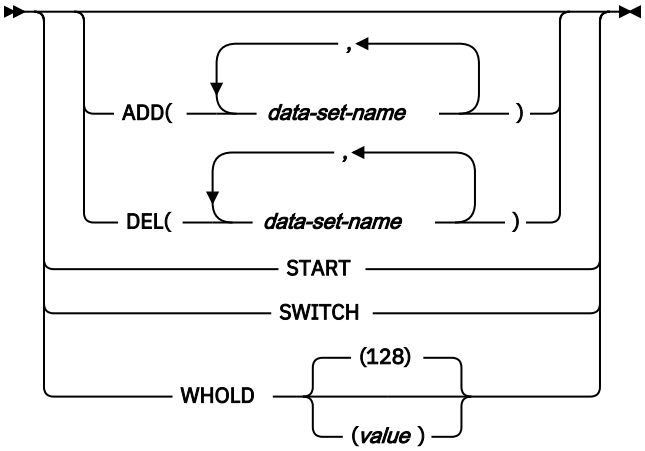

Specify at least one of the following suboptions:

- ADD|DEL
- START|STOP
- SWITCH|NOSWITCH
- WHOLD

ADD(data-set-name[,data-set-name])|DEL(data-set-name[,data-set-name])

Allows you to specify the name of the data set on which you want RMF to start or stop recording data. The name must match the name in the DEFINE CLUSTER statement. If you use a name that has not been defined, RMF issues a message.

ADD(data-set-name) allows RMF to use the specified data set to record sampled data. DEL(data-set-name) removes the specified data set from the list of data sets RMF uses to record data.

When you specify more than one data set name:

- Use a comma as a separator
- Specify no more than 100 data sets. If you specify more, RMF issues an error message
- Ensure that each data set name is unique

Examples:

- To specify two data sets for data set recording, use the following option:

```
DATASET(ADD(RMF.DS01))
DATASET(ADD(RMF.DS02))
```

RMF uses the empty data sets in the order in which they are defined. During data set recording, RMF writes the samples from its local storage buffer to the data sets. When all the data sets are full, RMF reuses the data sets, starting with the one that contains the oldest data.

- If you want to save data already recorded on a data set and make sure RMF does not reuse it, use the suboption DEL. This prevents RMF from writing over data in the specified data set. To save data contained in RMF.DS01, specified in the previous example, specify:

```
DATASET(DEL(RMF.DS01))
```

RMF does not reuse the data set during data set recording.

START|STOP
Allows you to start or stop data set recording. You can issue START|STOP at the beginning of a session on the session command START, or while the data gatherer is active with the session command MODIFY. If you do not want data set support for the data gatherer, use the default, which is DATASET(STOP).

RMF handles the START|STOP suboptions only at the end of a MINTIME. At this point, RMF has collected a set of samples representing the smallest sample time that the data reporter can display on the screen. By waiting until the end of the MINTIME to handle the START|STOP suboptions, RMF avoids recording partial sets of samples in the data sets.

SWITCH|NOSWITCH
Controls RMF's selection of a data set for recording sampled data.

If you specify SWITCH, RMF chooses the active data set as follows:

1. RMF searches for an empty data set to record samples
2. If there are no empty data sets, RMF reuses the data set with the oldest data

This option lets you reuse the specified data sets continuously, overlaying the oldest data once all the data sets are full.

If you specify NOSWITCH, or omit this suboption, RMF chooses the active data set as follows:

1. RMF searches for the data set with the most recent data and records samples if the data set is not full
2. If the data set with the most recent data is full, RMF searches for an empty data set to record samples
3. If there are no empty data sets, RMF reuses the data set with the oldest data

This option allows you to start the data gatherer and continue writing samples on a currently active data set that still has free space.

Note: NOSWITCH is effective only if specified or defaulted to when you start the data gatherer. It has no effect when specified on the session command MODIFY.

WHOLD(value)
Allows you to specify, in megabytes, a storage value that controls page releases in the RMF local storage buffer. The valid range of values for WHOLD is 1 to 999. RMF uses a default of 128 if you do not specify a value. If you specify a value outside the valid range, RMF uses 999 megabytes for a value above the range and 1 megabyte for a value below the range.

A page release discards the current and former copies of a page that are on central, expanded, or auxiliary storage, so that the page will not be read in before it is reused for new data. When the data in the local storage buffer has been copied to the data set and the storage amount exceeds the WHOLD value, the storage with duplicate data in the buffer becomes eligible for page release.

WHOLD works with the WSTOR option (see "WSTOR" on page 83) to control the page space needed for the storage buffer. You can specify a WHOLD value independent of the WSTOR value. If WHOLD is smaller than WSTOR:

- Page releases can occur before RMF uses all the storage in the local storage buffer

- When you turn data set recording off, the local storage buffer size assumes the WSTOR value.

If WHOLD is equal to or greater than WSTOR:

- Page releases occur once the WSTOR value is exceeded and RMF begins to wrap around the buffer.

When you activate data set recording, and the buffer contains data that the gatherer has already copied to the data set, the local storage buffer size reverts to the WHOLD value.

Starting data set support

Assume that before starting the data gatherer, you defined six VSAM data sets for data set recording. Issue the following START command to begin the data gatherer:

```
MODIFY  RMF,START III,MEMBER(08),DS(DEL(RMF.DS05),ADD(RMF.DS06),SWITCH)
```

You must identify the VSAM data set names to RMF through the DATASET option. The data set names must be identical to the names used to define the data sets, otherwise RMF will not recognize them.

Because MEMBER(08) is specified in the START command, RMF generates the member name ERBRMF08 and locates the member (normally found in SYS1.PARMLIB). Assume that ERBRMF08 contains the following DATASET options:

```
DATASET(START)
DATASET(ADD(RMF.DS01))
DATASET(ADD(RMF.DS02))
DATASET(ADD(RMF.DS03))
DATASET(ADD(RMF.DS04))
DATASET(ADD(RMF.DS05))
```

The default NOSWITCH at the beginning of this session permits RMF to continue writing on the active data set of the previous session (in this case, RMF.DS05).

Assume the following is true about the data sets at the beginning of this session:

- Data sets RMF.DS01 through RMF.DS04 are full
- RMF.DS05 is the active data set for this session
- RMF.DS06 is an empty data set.

With the DS options specified as parameters on the START session command, you modify the options as follows:

- Make a new data set available (ADD(RMF.DS06))
- Prevent RMF from writing on the currently active data set (DEL(RMF.DS05))
- Switch the recording of data to another data set (SWITCH).

START initiates data set recording, and RMF can use all the data sets listed with the ADD suboption.

As a result, RMF produces the following list of options following the rules of processing session options:

```
ERB305I  III : PARAMETERS
ERB305I  III : DATASET(WHOLD(128)) -- DEFAULT
ERB305I  III : DATASET(ADD(RMF.DS01)) -- MEMBER
ERB305I  III : DATASET(ADD(RMF.DS02)) -- MEMBER
ERB305I  III : DATASET(ADD(RMF.DS03)) -- MEMBER
ERB305I  III : DATASET(ADD(RMF.DS04)) -- MEMBER
ERB305I  III : DATASET(DEL(RMF.DS05)) -- COMMAND
ERB305I  III : DATASET(ADD(RMF.DS06)) -- COMMAND
ERB305I  III : DATASET(SWITCH) -- COMMAND
ERB305I  III : DATASET(START) -- MEMBER
ERB305I  III : MEMBER(08) -- COMMAND
ERB305I  III : WSTOR(128) -- DEFAULT
```

For more information, see Chapter 5, “How RMF processes session options,” on page 41.

RMF.DS06 is now available for data set recording. RMF.DS05 cannot be used for recording during the session. RMF.DS05 can be preallocated at the beginning of a TSO Monitor III reporter session and the

data on it displayed and analyzed. For more information, see Transferring Monitor III VSAM data sets to other systems in *z/OS RMF Reporter User's Guide*.

SWITCH causes RMF to switch to the next available data set, in this case, RMF.DS06 because it is empty. RMF.DS06 becomes the new active data set for this session. If you did not specify SWITCH in this example, data set recording would switch to an available data set anyway because RMF.DS05, the previously active data set, cannot be used. DATASET(DEL) has removed it from the list of data sets available for data set recording.

Note: If a data set contains the system ID or sysplex ID of another system or sysplex, Monitor III cannot overwrite this data set.

Modifying the data set support options

You can also modify the DATASET options while the data gatherer is active through the MODIFY session command. For more information, see "Modifying RMF session options" on page 37.

Example:

Assume you have started data set recording and have already defined data sets RMF.DS01 through RMF.DS05. Data sets RMF.DS01, RMF.DS02, RMF.DS03, and RMF.DS05 are full. RMF.DS01 contains the oldest data and RMF.DS04 is currently active.

You want to:

1. Save the data on RMF.DS04
2. Switch the current writing of the sampled data to another data set
3. Change the WHOLD value from the default of 128 to 256 megabytes.

The following command modifies the options:

```
MODIFY  RMF,MODIFY III,DS(SWITCH),DS(DEL(RMF.DS04)),DS(WHOLD(256))
```

1. The DEL suboption prevents RMF from overwriting data on RMF.DS04. RMF can no longer use RMF.DS04 for data set recording so the existing data is saved.
2. SWITCH causes RMF to begin writing in another data set. Because there is no empty data set, RMF chooses the data set with the oldest data, in this case RMF.DS01, and begins writing over the old data in it.
3. The WHOLD value lets RMF hold a copy in its buffer of 256 megabytes of storage containing data already copied to the data set. After it exceeds the value, it begins to page release the storage in the buffer containing the duplicate data.

Stopping data set support

You can stop the data gatherer from writing to any data set or never activate data set recording. If you do not want the data set support for a data gatherer session, you can do one of the following:

- Specify the DATASET(STOP) option in the PARM field of the START session command
- Specify the DATASET(STOP) option in the PARM field of the MODIFY session command
- Specify the DATASET(STOP) option in an RMF Parmlib member
- Use the default DATASET(STOP).

You can also use the DATASET(STOP) option to suspend recording until you need it. You can activate recording by overriding DATASET(STOP) with DATASET(START) on a session START or MODIFY command.

Example:

Parmlib member ERBRMF04 may contain the following:

```
DATASET(STOP)
DATASET(ADD(RMF.DS01))
DATASET(ADD(RMF.DS02))
```

```
DATASET(ADD(RMF.DS03))
DATASET(ADD(RMF.DS04))
DATASET(ADD(RMF.DS05))
```

The DS(STOP) in the member means that no active data set recording occurs when a data gatherer session is started.

To start data set recording later, specify:

```
F RMF,S III,DS(START)
```

or

```
F RMF,F III,DS(START)
```

The DS(START) option on the command overrides the DS(STOP) option in Parmlib member ERBRMF04, and permits the recording of sampled data to the data sets defined by the DS(ADD) options.

If you want, you can also change the data set names specified in the DS(ADD) options.

Data set support for daylight saving time

Data set support works as follows when local time is changed:

- Time is set forth (winter to summer time):

  There is a gap in local time where no data is selected. When the currently active data set is full, the data set with the oldest data is selected to store the current data.
- Time is set back (summer to winter time):

  The data on the data set(s) with a time stamp of the future is deleted and recording on the data set continues.

  Note: When time set back, there is a time window where data is collected twice with the same local time stamp. The existing data is deleted. When the existing data in the overlapping time window is essential for your monitoring, you may remove this data set(s) from RMFGAT (DS(DEL(name)) before time change. These data is archived now and can be used with the Monitor III reporter when allocated as RMFDS00 at a TSO session. For more details, see Data set allocation in *z/OS RMF Reporter User's Guide*.

Appendix A. Accessibility

Accessible publications for this product are offered through IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSLTBW/welcome).

If you experience difficulty with the accessibility of any z/OS information, send a detailed message to the Contact the z/OS team web page (www.ibm.com/systems/campaignmail/z/zos/contact_z) or use the following mailing address.

IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Building 707
2455 South Road
Poughkeepsie, NY 12601-5400
United States

Accessibility features

Accessibility features help users who have physical disabilities such as restricted mobility or limited vision use software products successfully. The accessibility features in z/OS can help users do the following tasks:

- Run assistive technology such as screen readers and screen magnifier software.
- Operate specific or equivalent features by using the keyboard.
- Customize display attributes such as color, contrast, and font size.

Consult assistive technologies

Assistive technology products such as screen readers function with the user interfaces found in z/OS. Consult the product information for the specific assistive technology product that is used to access z/OS interfaces.

Keyboard navigation of the user interface

You can access z/OS user interfaces with TSO/E or ISPF. The following information describes how to use TSO/E and ISPF, including the use of keyboard shortcuts and function keys (PF keys). Each guide includes the default settings for the PF keys.

- *z/OS TSO/E Primer*
- *z/OS TSO/E User's Guide*
- *z/OS ISPF User's Guide Vol I*

Dotted decimal syntax diagrams

Syntax diagrams are provided in dotted decimal format for users who access IBM Knowledge Center with a screen reader. In dotted decimal format, each syntax element is written on a separate line. If two or more syntax elements are always present together (or always absent together), they can appear on the same line because they are considered a single compound syntax element.

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that the screen reader is set to read out punctuation. All the syntax elements that have the same dotted decimal number (for example, all the syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, your syntax can include either USERID or SYSTEMID, but not both.

The dotted decimal numbering level denotes the level of nesting. For example, if a syntax element with dotted decimal number 3 is followed by a series of syntax elements with dotted decimal number 3.1, all the syntax elements numbered 3.1 are subordinate to the syntax element numbered 3.

Certain words and symbols are used next to the dotted decimal numbers to add information about the syntax elements. Occasionally, these words and symbols might occur at the beginning of the element itself. For ease of identification, if the word or symbol is a part of the syntax element, it is preceded by the backslash (\) character. The * symbol is placed next to a dotted decimal number to indicate that the syntax element repeats. For example, syntax element *FILE with dotted decimal number 3 is given the format 3 * FILE. Format 3* FILE indicates that syntax element FILE repeats. Format 3* * FILE indicates that syntax element * FILE repeats.

Characters such as commas, which are used to separate a string of syntax elements, are shown in the syntax just before the items they separate. These characters can appear on the same line as each item, or on a separate line with the same dotted decimal number as the relevant items. The line can also show another symbol to provide information about the syntax elements. For example, the lines 5.1*, 5.1 LASTRUN, and 5.1 DELETE mean that if you use more than one of the LASTRUN and DELETE syntax elements, the elements must be separated by a comma. If no separator is given, assume that you use a blank to separate each syntax element.

If a syntax element is preceded by the % symbol, it indicates a reference that is defined elsewhere. The string that follows the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you must refer to separate syntax fragment OP1.

The following symbols are used next to the dotted decimal numbers.

? indicates an optional syntax element
The question mark (?) symbol indicates an optional syntax element. A dotted decimal number followed by the question mark symbol (?) indicates that all the syntax elements with a corresponding dotted decimal number, and any subordinate syntax elements, are optional. If there is only one syntax element with a dotted decimal number, the ? symbol is displayed on the same line as the syntax element, (for example 5? NOTIFY). If there is more than one syntax element with a dotted decimal number, the ? symbol is displayed on a line by itself, followed by the syntax elements that are optional. For example, if you hear the lines 5 ?, 5 NOTIFY, and 5 UPDATE, you know that the syntax elements NOTIFY and UPDATE are optional. That is, you can choose one or none of them. The ? symbol is equivalent to a bypass line in a railroad diagram.

! indicates a default syntax element
The exclamation mark (!) symbol indicates a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicate that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the dotted decimal number can specify the ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the default option for the FILE keyword. In the example, if you include the FILE keyword, but do not specify an option, the default option KEEP is applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, the default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP applies only to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.

*** indicates an optional syntax element that is repeatable**
The asterisk or glyph (*) symbol indicates a syntax element that can be repeated zero or more times. A dotted decimal number followed by the * symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1* data area, you know that you can include one data area, more than one data area, or no data area.

If you hear the lines `3* , 3 HOST, 3 STATE`, you know that you can include `HOST, STATE`, both together, or nothing.

Notes:

1. If a dotted decimal number has an asterisk (*) next to it and there is only one item with that dotted decimal number, you can repeat that same item more than once.
2. If a dotted decimal number has an asterisk next to it and several items have that dotted decimal number, you can use more than one item from the list, but you cannot use the items more than once each. In the previous example, you can write `HOST STATE`, but you cannot write `HOST HOST`.
3. The * symbol is equivalent to a loopback line in a railroad syntax diagram.

+ indicates a syntax element that must be included

The plus (+) symbol indicates a syntax element that must be included at least once. A dotted decimal number followed by the + symbol indicates that the syntax element must be included one or more times. That is, it must be included at least once and can be repeated. For example, if you hear the line `6.1+ data area`, you must include at least one data area. If you hear the lines `2+, 2 HOST,` and `2 STATE`, you know that you must include `HOST, STATE,` or both. Similar to the * symbol, the + symbol can repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loopback line in a railroad syntax diagram.

Notices

This information was developed for products and services that are offered in the USA or elsewhere.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
United States of America

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

This information could include missing, incorrect, or broken hyperlinks. Hyperlinks are maintained in only the HTML plug-in output for the Knowledge Centers. Use of hyperlinks in other output formats of this information is at your own risk.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Site Counsel
2455 South Road

Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or

reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's name, email address, phone number, or other personally identifiable information for purposes of enhanced user usability and single sign-on configuration. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at ibm.com/privacy and IBM's Online Privacy Statement at ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at ibm.com/software/info/product-privacy.

Policy for unsupported hardware

Various z/OS elements, such as DFSMS, JES2, JES3, and MVS, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

Minimum supported hardware

The minimum supported hardware for z/OS releases identified in z/OS announcements can subsequently change when service for particular servers or devices is withdrawn. Likewise, the levels of other software products supported on a particular release of z/OS are subject to the service support lifecycle of those products. Therefore, z/OS and its product publications (for example, panels, samples, messages, and product documentation) can include references to hardware and software that is no longer supported.

- For information about software support lifecycle, see: IBM Lifecycle Support for z/OS (www.ibm.com/software/support/systemsz/lifecycle)
- For information about currently-supported IBM hardware, contact your IBM representative.

Programming Interface Information

This book is intended to help the customer to use RMF sessions. It contains a description of what RMF is, what it can do, and how to use the different sessions.

The book also documents intended Programming Interfaces that allow the customer to write programs to obtain the services of RMF.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and Trademark information (www.ibm.com/legal/copytrade.shtml).

Java™ is a trademark of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Windows, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.

Glossary

This glossary contains chiefly definitions of terms used in this book, but some more general RMF and MVS terms are also defined.

Words that are set in *italics* in the definitions are terms that are themselves defined in the glossary.

APPC/MVS
Advanced program-to-program communication

ASCH address space
APPC transaction scheduler address space

AS
Address space

address space
That part of MVS main storage that is allocated to a job.

auxiliary storage (AUX)
All addressable storage, other than main storage, that can be accessed by means of an I/O channel; for example storage on direct access devices.

background session
In RMF, a monitor session that is started and controlled from the operator console. Contrast with *interactive session*

balanced systems
To avoid bottlenecks, the system resources (CP, I/O, storage) need to be balanced.

basic mode
A central processor mode that does not use logical partitioning. Contrast with *logically partitioned (LPAR) mode*.

bottleneck
A system resource that is unable to process work at the rate it comes in, thus creating a queue.

callable services
Parts of a program product that have a published external interface and can be used by application programs to interact with the product.

captured storage
See shared page group.

capture ratio
The ratio of reported CPU time to total used CPU time.

central processor (CP)
The part of the computer that contains the sequencing and processing facilities for instruction execution, initial program load, and other machine operations.

central processor complex (CPC)
A physical collection of hardware that consists of central storage, one or more central processors, timers, and channels.

channel path
The channel path is the physical interface that connects control units and devices to the CPU.

CICS
Customer Information Control System

CIM provider
A CIM provider is the link between the CIM server and the system interfaces. It allows the CIM server to access and manage the resources. Each CIM provider exposes the resources it represents in a standard way, using a small number of classes from the CIM schema or derived from the CIM schema. RMF monitoring providers are CIM providers implemented by RMF.

contention
Two or more incompatible requests for the same resource. For example, contention occurs if a user requests a resource and specifies exclusive use, and another user requests the same resource, but specifies shared use.

coupling facility
See *Cross-system Extended Services/Coupling Facility*.

CP
Central processor

criteria
Performance criteria set in the WFEX report options. You can set criteria for all report classes (PROC, SYSTEM, TSO, and so on).

CPU speed
Measurement of how much work your CPU can do in a certain amount of time.

cross-system coupling facility (XCF)
A component of MVS that provides functions to support cooperation between authorized programs running within a *sysplex*.

Cross-system Extended Services/Coupling Facility (XES/CF)
Provides services for MVS systems in a sysplex to share data on a coupling facility (CF).

CS
Central storage

Customer Information Control System (CICS)
An IBM licensed program that enables transactions entered at remote terminals to be processed concurrently by user-written application programs. It includes facilities for building, using, and maintaining data bases.

cycle
In RMF, the time at the end of which one sample is taken. Varies between 50 ms and 9999 ms. See also *sample*.

data sample
See *sample*

DCM
See *Dynamic Channel Path Management*

delay
The delay of an address space represents a job that needs one or more resources but that must wait because it is contending for the resource(s) with other users in the system.

direct access storage device (DASD)
A device in which the access time is effectively independent of the location of the data. Usually: a magnetic disk device.

DLY
Delay

DP
Dispatching priority

dynamic channel path management
Dynamic channel path management provides the capability to dynamically assign channels to control units in order to respond to peaks in demand for I/O channel bandwidth. This is possible by allowing you to define pools of so-called floating channels that are not related to a specific control unit. With the help of the Workload Manager, channels can float between control units to best service the work according to their goals and their importance.

EMIF
ESCON multiple image facility

enclave
An enclave is a group of associated dispatchable units. More specifically, an enclave is a group of SRB routines that are to be managed and reported on as an entity.

EPDM
Enterprise Performance Data Manager/MVS

execution velocity
A measure of how fast work should run when ready, without being delayed for processor or storage access.

exception reporting
In RMF, the reporting of performance measurements that do not meet user-defined criteria. Shows potential performance problems explicitly, thus avoiding the need for constant monitoring.

generalized trace facility (GTF)
A service program that records significant system events, such as supervisor calls and start I/O operations, for the purpose of problem determination.

GO mode
In RMF, the Monitor III mode in which the screen is updated with the interval you specified in your session options. The terminal cannot be used for anything else when it is in GO mode. See also *mode*.

graphic mode
In RMF Monitor III, the mode which presents the performance data from the system in graphic format using the GDDM product. Contrast with *tabular mode*.

GTF
generalized trace facility

high-speed buffer (HSB)
A cache or a set of logically partitioned blocks that provides significantly faster access to instructions and data than provided by central storage.

HS
hiperspace

HSB
High-speed buffer

HSM
Hierarchical Storage Manager

IBM Z Application Assist Processor (zAAP)
A special purpose processor configured for running Java programming on selected zSeries machines.

IBM Z Integrated Information Processor (zIIP)
A special purpose processor designed to help free-up general computing capacity and lower overall total cost of computing for selected data and transaction processing workloads for business intelligence (BI), ERP and CRM, and selected network encryption workloads on the mainframe.

IMS
Information Management System

Information Management System (IMS)
A database/data communication (DB/DC) system that can manage complex databases and networks. Synonymous with IMS/VS.

interactive session
In RMF, a monitor display-session that is controlled from the display terminal. Contrast with *background session*.

JES
Job Entry Subsystem

LCU
Logical control unit. Logical control units are also called 'Control Unit Headers ' (CUH). For details about LCU/CUH please refer to the applicable *IBM Z Input/Output Configuration Program User's Guide for ICP IOCP* (SB10-7037).

logically partitioned (LPAR) mode
A central processor mode that is available on the Configuration frame when using the PR/SM feature. It allows an operator to allocate processor unit hardware resources among logical partitions. Contrast with *basic mode*.

logical partition (LP)
A subset of the processor hardware that is defined to support an operating system. See also *logically partitioned (LPAR) mode*.

LP
Logical partition

LPAR
Logically partitioned (mode)

LPAR cluster
An LPAR cluster is the subset of the systems that are running as LPARs on the same CEC. Based on business goals, WLM can direct PR/SM to enable or disable CP capacity for an LPAR, without human intervention.

migration rate
The rate (pages/second) of pages being moved from expanded storage through central storage to auxiliary storage.

mintime
The smallest unit of sampling in Monitor III. Specifies a time interval during which the system is sampled. The data gatherer combines all samples gathered into a set of samples. The set of samples can be summarized and reported by the reporter.

mode
Monitor III can run in various modes: GO mode (see *GO mode*) and STOP mode, which is the default mode. See also *graphic mode* and *tabular mode*.

MPL
Multiprogramming level

OMVS
Reference to z/OS UNIX System Services

partitioned data set (PDS)
A data set in direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

PDS
partitioned data set

performance management
The activity which monitors and allocates data processing resources to applications according to goals defined in a service level agreement or other objectives.

The discipline that encompasses collection of performance data and tuning of resources.

PR/SM
Processor Resource/Systems Manager

Processor Resource/Systems Manager (PR/SM)
The feature that allows the processor to run several operating systems environments simultaneously and provides logical partitioning capability. See also *LPAR*.

range
The time interval you choose for your report.

Resident time
The time the address space was swapped in, in units of seconds.

RMF monitoring provider
see CIM provider

sample
Once in every cycle, the number of jobs waiting for a resource, and what job is using the resource at that moment, are gathered for all resources of a system by Monitor III. These numbers constitute one sample.

SCP
System control program

seek
The DASD arm movement to a cylinder. A seek can range from the minimum to the maximum seek time of a device. In addition, some I/O operations involve multiple imbedded seeks where the total seek time can be more than the maximum device seek time.

service class
In Workload Manager, a subdivision of a *workload*. Performance goals and capacity boundaries are assigned to service classes.

service level agreement (SLA)
A written agreement of the information systems (I/S) service to be provided to the users of a computing installation.

Service Level Reporter (SLR)
An IBM licensed program that provides the user with a coordinated set of tools and techniques and consistent information to help manage the data processing installation. For example, SLR extracts information from SMF, IMS, and CICS logs, formats selected information into tabular or graphic reports, and gives assistance in maintaining database tables.

service rate
In the system resources manager, a measure of the rate at which system resources (services) are provided to individual jobs. It is used by the installation to specify performance objectives, and used by the workload manager to track the progress of individual jobs. Service is a linear combination of processing unit, I/O, and main storage measures that can be adjusted by the installation.

shared page groups
An address space can decide to share its storage with other address spaces using a function of RSM. As soon as other address spaces use these storage areas, they can no longer be tied to only one address space. These storage areas then reside as *shared page groups* in the system. The pages of shared page groups can reside in central, expanded, or auxiliary storage.

SLA
service level agreement

SLIP
serviceability level indication processing

SLR
Service Level Reporter

SMF
System management facility

SMF buffer
A wrap-around buffer area in storage, to which RMF data gatherers write performance data, and from which the Postprocessor extracts data for reports.

speed
See *workflow*

SRB
Service request block

SRM
System resource manager

SSCH
Start subchannel

system control program (SCP)
Programming that is fundamental to the operation of the system. SCPs include MVS, VM, and VSE operating systems and any other programming that is used to operate and maintain the system. Synonymous with *operating system*.

sysplex
A complex consisting of a number of coupled MVS systems.

tabular mode
In RMF, the mode in which Monitor III displays performance data in the form of lists. Contrast with *graphic mode*.

TCB
Task control block

threshold
The exception criteria defined on the report options screen.

throughput
A measure of the amount of work performed by a computer system over a period of time, for example, number of jobs per day.

TPNS
Teleprocessing network simulator

TSO
Time Sharing Option, see *Time Sharing Option/Extensions*

Time Sharing Option Extensions (TSO/E)
In MVS, a time-sharing system accessed from a terminal that allows user access to MVS system services and interactive facilities.

UIC
Unreferenced interval count

uncaptured time
CPU time not allocated to a specific address space.

using
Jobs getting service from hardware resources (PROC or DEV) are *using* these resources.

velocity
A measure of how fast work should run when ready, without being delayed for processor or storage access. See also *execution velocity*.

VTOC
Volume table of contents

workflow
The workflow of an address space represents how a job uses system resources and the speed at which the job moves through the system in relation to the maximum average speed at which the job could move through the system.

The workflow of resources indicates how efficiently users are being served.

workload
A logical group of work to be tracked, managed, and reported as a unit. Also, a logical group of service classes.

WLM
Workload Manager

XCF
Cross-system coupling facility

XES/CF
See *Cross-system Extended Services/Coupling Facility*.

zAAP
see IBM Z Application Assist Processor.

zIIP
see IBM Z Integrated Information Processor.

Index

D

E

F

G

H

I

J

K

L

M

N

O

P

R

S

T

U

V

W

Z

Product Number: 5650-ZOS

SC27-4934-40